

CA Spectrum®

Modeling and Managing Your IT Infrastructure Administrator Guide

Release 9.3



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

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

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

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

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

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

The manufacturer of this Documentation is CA.

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

Copyright © 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 Spectrum®
- CA Spectrum® Modeling Gateway Toolkit (Modeling Gateway)
- CA Spectrum® Southbound Gateway Toolkit (Southbound Gateway)
- CA Spectrum® Network Configuration Manager (NCM)
- CA Spectrum® Report Manager
- CA Spectrum® Service Manager

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: Getting Started	13
Network Modeling in CA Spectrum.....	13
OneClick Topologies	13
Universe Topology	14
Global Collections Topology.....	16
World Topology.....	17
TopOrg Topology.....	18
Topology Toolbar	19
Icons in Topology Views	20
Aggregate Icons.....	21
Individual Icons	21
Icon Types by Theme.....	22
Icon Color and Condition.....	24
Provision Access to Modeled Elements	26
 Chapter 2: Discovering and Modeling Your Network	 27
Methods for Modeling Network Entities	27
Discovery	28
Separate Configurations.....	29
Discovery Console	30
Configuration Tab.....	30
Discovery Tab	33
Modeling Tab	36
History Tab	38
Discovery Connection Status.....	40
Open the Discovery Console	40
Define a Discovery Configuration.....	41
How to Set Discovery Configuration Parameters.....	43
Seed Routers	43
IP/Host Name Boundary List	44
SNMP Information	46
Modeling Options.....	48
Advanced Options	54
Scheduling Options	56
Activate a Discovery Session	56
Activate a Modeling Session	57

Run a Network Services Discovery	58
Create Discovery Configuration Folders.....	59
Reorganize Discovery Configurations.....	59
Rename Discovery Configurations or Folders	60
VLAN Discovery	60
View, Filter, and Export Results Lists.....	60
Export a Results List	61
Set a Modeling Results List to Export Automatically	61
Filter Results Using Advanced Filter	62
Import Discovery Configurations	64
Export Discovery Configurations	64
After Discovering and Modeling.....	65
VNM AutoDiscovery Control Settings	65
Access VNM AutoDiscovery Control	66
Loopback Interfaces and Discovery.....	66

Chapter 3: Modeling Your Network Manually 67

When to Model Manually in OneClick	67
How to Model Manually in the Universe Topology.....	68
Create Model Dialog	68
Add Containers to Universe Topology Views	72
Add Network Devices to Universe Topology Views	73
Remove Modeled Elements from the Universe Topology View	85
Delete Modeled Elements from the Universe Topology View	86
Cut Modeled Elements from the Topology View or List View	86
Enhance Topology Views	86
Export a Topology View	87
Modeling Manually in a Global Collections Topology	87
Dynamic Membership	87
Static Membership	88
Connections Between Modeled Elements (Members)	88
Updating Modeled Elements in Global Collections.....	88
Generate Reports on Collections	89
How to Define and Manage Global Collections	89
Modeling Manually in the World Topology View.....	106
How to Model Locations	107
Define a Top-Level or Sub-Level Location View	108
Modeling Manually in the TopOrg Topology View.....	109
How to Model Services in the TopOrg Topology.....	109
Using Favorites	111
Deleting from Favorites.....	112

Lost and Found Model Information Subview	112
--	-----

Chapter 4: Configuring Modeled Devices and Interfaces 115

Device and Interface Threshold Settings	115
Device Threshold Settings	115
Interface Threshold Settings	119
Update Device Interface and Connection Information	119
Automatic Updates of Device Interface and Connection Information	120
Manually Updating Interface and Connection Information	122
Access Interface and Connection Update Controls	126
Redundant Connections Between CA Spectrum and Modeled Devices	128
Redundancy Preferred Addresses List	128
IP Redundancy Subview	129
Select Preferred Redundant Addresses	129
Exclude Redundant Addresses	130
Interface Reconfigurations	131
Primary IP Address Modification	131
Change the Primary IP Address for a Device in the Preferred Address List	132
Change the Device IP Address to an Interface Address	132
Change the Primary IP Address for a Device to use an Interface Secondary IP Address	133
IPv6 Information	134

Chapter 5: Editing and Enhancing Topology Views 135

Topology Edit Mode	135
Access Edit Mode	135
Edit Mode Toolbar	136
Set Topology View Edit Mode Preferences	137
Modifying Topology Views	138
Multi-User Considerations	139
Resize Model Icons	139
Add Shapes, Lines, or Text to a View	140
Change Shapes, Lines, and Text Characteristics	141
Background Editor	142
Modify the Topology Background	143
Group Items in a View	144
Ungroup Items in a View	145
Send Items to the Back	145
Bring Items to the Front	145

Chapter 6: Modifying Model Attributes 147

Model Attributes	147
Attributes in the Information Tab	148
VNM Attributes in the Information Tab	149
General Information Subview	149
CA Spectrum Modeling Information Subview	149
Online Database Backup Subview	149
SpectroSERVER Control Subview	150
How Trap Storm Detection Works	156
AutoDiscovery Control Subview	157
Fault Isolation Subview	161
Live Pipes Subview	161
Alarm Management Subview	162
BGP Manager Subview	164
Network Configuration Manager Subview	165
Thresholds And Watches Subview	166
Host Security Information Subview	167
Modeling Gateway Subview	167
IP Services Subview	167
Logical Connection Import Subview	167
Shared IP Detection and Alarming	168
Attributes Tab	170
Access Attributes from the Attributes Tab	171
Edit Attributes in the Attributes Tab	171
Edit Multiple Attributes at Once in the Attributes Tab	172
Examine the Same Attribute on Multiple Models	173
View List Attribute Values	173
Update Attribute Values	174
OneClick Attribute Editor	174
Open Attribute Editor	175
Attribute Editor Dialog	175
Attribute Edit Results Dialog	177
User-Defined Attributes	178
Create User-Defined Attributes	178
Change Attributes in Conjunction with Search	178
Example: Define a Search to Create an Attribute for Editing	179
Edit Attributes for Specific Devices or for Model Types	180
Model Type Reevaluation	183
Edit the Model Type Reevaluation Interval	183
Change Management Attributes	184
Interface Configuration Attributes	185

Stale Interfaces	185
Maintenance Mode Attributes.....	186
Rollup Alarm Attributes.....	186
Model Status and Alarm Conditions	188
Rollup Condition Thresholds	188
SNMP Communication Attributes	189
Threshold Attributes	190
How CA Spectrum Calculates CPU and Memory Utilization	191
Normalized CPU Utilization Calculation Requirements	192
Normalized Memory Utilization Calculation Requirements	193
Normalized CPU Utilization Attributes.....	193
Normalized Memory Utilization Attributes.....	195
Calculate Normalized CPU Utilization	197
Calculate Normalized Memory Utilization	198
Troubleshoot CPU and Memory Utilization Calculation	200

Chapter 7: Fault Management 201

Fault Isolation Settings	201
Port Fault Correlation.....	203
Port Fault Correlation Options	204
Port Fault Correlation Criteria	205
Port Fault Correlation Caveats	205
Example: Port Fault Correlation Scenario 1	206
Example: Port Fault Correlation Scenario 2	207
Example: Port Fault Correlation Scenario 3	210
Port Fault Correlation Anomalies.....	211
Configure Cross-Landscape Fault Correlation	212
Designate a Model as a Proxy Model.....	212
Cross-Landscape Fault Correlation Example	213
Configuring Port Status Monitoring	214
Port Status Polling Criteria	216
Port Status Events and Alarms	217
Link Traps	218
Interface Trap Configuration.....	219
Wide Area Link Monitoring	221
LinkFaultDisposition	221
Wide Area Link Monitoring Scenarios.....	222
Wide Area Link Modeling Best Practices.....	223
Port Layer Alarm Suppression	224
Port Criticality.....	224
Live Pipes and Fault Management	224

Enable or Disable Live Pipes System-Wide.....	225
Enable or Disable Live Pipes on Individual Links	225
Receiving Port Alarms	225
Monitoring Physical and Logical Connections.....	227
Suggested Port Fault Settings for Optimal Fault Notification	229
Device Criticality	230
Configuring Fault Management for Pingables.....	230
Connect Pingable Models	231
Mapping Traps from Other Models to Pingable Models	231
False Management Lost or Contact Lost Alarms	233

Chapter 8: Modeling and Managing SNMPv3 Devices 235

SNMPv3 Support	235
SNMPv3 Authentication	235
Enable SNMPv3 Privacy.....	236
64-Bit Counters	238
SNMPv3 Support Issues	239
Edit SNMP v3 Profiles Dialog.....	239
Manually Model an SNMPv3 Device	240
Modeling an SNMPv3 Device Using a CA Spectrum Toolkit.....	243
Model an SNMP v2c Device Using a CA Spectrum Toolkit	245
Change Security Information for a Device Model	245
Add Context Name Information.....	246
Specify an Authentication Encryption Algorithm on a Per-Model Basis	246
Specify a Privacy Encryption Algorithm on a Per-Model Basis.....	249
Troubleshoot SNMPv3 Communication Issues	251

Chapter 9: CA Spectrum Intelligence 253

Inductive Modeling Technology	253
Static Configuration of Device Models.....	253
Dynamic Configuration of Device Models.....	254
Pulled Board List.....	254
Router Reconfiguration Events	255
Condition Versus Rollup Condition	256
Attributes Determining Condition and Rollup Condition	256
Condition and Rollup Condition Sensitivity.....	258
Rollup Condition Flow	260
Example of Rollup Condition Propagation	261
Example Rollup Condition Process.....	263
Fault Isolation.....	264
How Model Category Affects Contact Status.....	265

Fault Isolation Examples	268
Duplicate Addresses	272
Manually Clear Duplicate Addresses	274
Automatic Naming and Addressing	274
Detection of Firmware Problems	275
Interface Intelligence	275
Interface Alarms	277
Interface Events	278
 Glossary	 281
 Index	 285

Chapter 1: Getting Started

This section contains the following topics:

[Network Modeling in CA Spectrum](#) (see page 13)

[OneClick Topologies](#) (see page 13)

[Icons in Topology Views](#) (see page 20)

[Provision Access to Modeled Elements](#) (see page 26)

Network Modeling in CA Spectrum

Network modeling in CA Spectrum is the act of graphically representing network entities and their connections. Icons that are created, placed, and connected within the OneClick topology views represent various aspects of a modeled network.

Using the modeling features of the OneClick client, you can easily create and maintain accurate software models of your network. These intelligent network models enable CA Spectrum to determine actual points of failure and to suppress superfluous alarms.

CA Spectrum network representation is based on logical relationships and rules and appears different from your network diagrams. Discovery uses address tables and ICMP ping tests to identify subnet address ranges and devices within those ranges. Once discovered, CA Spectrum models those devices and subnets.

OneClick Topologies

You can use four core topologies to model your IT infrastructure in CA Spectrum:

- [Universe Topology](#) (see page 14)
- [Global Collections Topology](#) (see page 16)
- [World Topology](#) (see page 17)
- [TopOrg Topology](#) (see page 18)

All four of these topologies are available from the Navigation panel.

Note: We recommend that you begin modeling with the Universe topology. After you have established one or more modeled elements in the Universe topology, you can reuse these modeled elements to define the other topologies.

To navigate through the model views of any topology, click the view control icons in the toolbar. In some cases, you can click an [aggregate icon container](#) (see page 20) to view its content.

Universe Topology


The Universe topology helps you organize an enterprise network view of your infrastructure. Most often, it provides a view of:

- A top-level topology view of OSI Layer 3 devices and their connections
- A drill-down topology view of OSI Layer 2 devices and their connections
- A Component Detail view, showing the attributes that are associated with a modeled entity


The top view in the Universe topology typically includes:

- The CA Spectrum Virtual Network Machine (VNM)
- Network groupings
- Network segments
- OSI Layer 3 devices and their connections

The following illustration shows a typical top view:




The following illustration shows a drill-down view of a LAN container that is selected from the top view. A drill-down view in the Universe topology most often includes all OSI Layer 2 devices and their connections. The drill-down view also shows off-page references to devices modeled in other views, as shown in the following illustration:



Component Detail Panel

The Component Detail panel within the Universe topology identifies the attributes that are associated with a modeled network entity. Attributes can include the device interfaces, alarms and events, and other pertinent device information.



You can view the device attributes and possibly change their settings by clicking the Component Detail panel tabs. Depending on the context of the Contents panel, you can use the Component Detail panel to:

- View current alarms in the Alarm Details tab.
- View and modify general device settings in the Information tab. For example, grant or deny access to a modeled device by providing or possibly removing a security string.
- View root cause analysis data in the Root Cause tab.
- View CPU and memory utilization information in the Performance tab.
- View device interface information in the Interfaces tab.
- View neighboring routers in the Neighbors tab.
- View historical events in the Events tab.
- View attribute information from the Attributes tab. This information appears only under these conditions:
 - You select an entity from the Explorer tab.
 - You are in either the Topology, List, or Events tab in the Contents panel.

Define Models in the Universe Topology

You can define models in the Universe topology using the OneClick Discovery feature, which automates the modeling process for you. You can also manually define new models or edit existing models in the Universe topology by using the modeling tools that are provided with OneClick.

The Universe topology view represents a true connectivity view of your infrastructure. Therefore, we strongly recommend that you reuse modeled elements from this view when creating other views. Therefore, as a best practice, copy model elements from the Universe topology view to create Global Collections, World, or TopOrg views. This approach helps to ensure accurate fault isolation of your network within the OneClick environment.

Global Collections Topology

A *landscape* is the network domain that is managed by a single SpectroSERVER. In OneClick, a landscape is the network view of one SpectroSERVER. To organize entity-based network views that span one or more landscapes, use Global Collections. Global collections enable you to monitor all aspects of your IT infrastructure from any perspective.

As an administrator, you can use Global Collections to create and track collections of network entities, organizations, or services that make up your infrastructure. For example, you can create and maintain collections that identify and track:

- Response teams within an organization responsible for maintaining equipment
- Devices supporting various services in your organization
- Customers receiving services from your organization

More information:

[Modeling Manually in a Global Collections Topology](#) (see page 87)


View or Modify Modeled Devices

You can view and change modeled device attributes or settings in the Component Detail panel within a Global Collections topology. Click the tabs in this panel for general information about a device, its interfaces, alarms, events, and other pertinent information.

World Topology

The World topology helps organize your network geographically in OneClick. In this topology, you can represent device models of network locations from a national or regional level all the way down to a wiring closet.

The following example illustrates a drill-down view of an equipment room that is at a fictitious North Dallas location.



In the World topology, you can create several layers of views that represent locations of your network devices. For example, you can have views for Texas regional offices, Dallas office, and Dallas Equipment room. Additionally, the Component Detail panel lets you view and sometimes change the attributes that are associated with a modeled device in any World topology view. For instance, clicking the Component Detail tabs for a modeled device lets you view device information, interfaces, alarms, events, and other pertinent device information.

Note: When populating the World topology views with modeled devices, we highly recommend that you copy and paste modeled elements from Universe topology views. Universe topology views represent the true connectivity views of your infrastructure, helping to ensure accurate fault isolation of your network within the OneClick environment.


More information:

[Modeling Manually in the World Topology View](#) (see page 106)

TopOrg Topology

The TopOrg topology allows you to represent your network organizationally. In this topology, you can group subnets and device models by services, responsibilities, departments, or by other organizational considerations.

The following illustration displays an organizational view that identifies individuals and groups within a fictitious Finance Department. This type of view is useful when identifying how a network failure or a reconfiguration affects an organizational unit.



If you purchased the Service Manager module, you can use this module with the TopOrg topology to model business services and applications. Further, the Service Manager module also tracks the performance of the service against a contract or Service Level Agreement (SLA).

Note: For more information about using the Service Manager module, see the *Service Manager User Guide*.

In the TopOrg topology, you can create several layers of views that represent various levels of your network devices. For example, you can have views for Enterprise ownership, Department ownership, supporting devices, and supporting services. Additionally, the Component Detail panel lets you view and possibly change the attributes that are associated with a modeled device in any TopOrg topology view.

Note: When populating the TopOrg topology views with modeled devices, we recommend that you copy and paste modeled elements from Universe topology views. Universe topology views represent the true connectivity views of your infrastructure. Therefore, using them as your base for all other views helps ensure more accurate fault isolation of your network within the OneClick environment.

More information:

[Modeling Manually in the TopOrg Topology View](#) (see page 109)

Topology Toolbar

The following table describes some of the buttons available in the Topology tab toolbar for working with topologies.

Icon	Description
	Edit mode: Click to put the current topology view into Edit mode.
	Create new model by type: Click to create a new model by type and add it to the topology view.
	Create new model by IP: Click to create a new model by IP address or Host Name. Click the down arrow and select one of the following options: ■ Create By IP ■ Create By Host Name
	Discovery: Click to create a new discovery based on the selected model or models.

Icon	Description
	Spotlight View: Click to highlight all models related to a VPN, a VLAN, or router redundancy in the Topology view. Spotlighting allows you to easily determine relationships between these items and your network, and between these items and other models on your network. When you click the Spotlight view button, a menu appears containing the following options: ■ Router Redundancy ■ VLAN List ■ VPN List
	Remove Model: Click to remove the selected model from the Topology view.

More information:

[Topology Edit Mode](#) (see page 135)

Icons in Topology Views

Icons that appear in OneClick topology views are graphical representations of network entities. Some network entities include:

- Individual devices
- Groups of devices
- Geographic locations
- Physical connections

An icon is simply the image with which you interact when you manipulate and configure a modeled element. When monitoring the condition of your network, an icon represents the current status of a device, network group, device location, or a physical link.

CA Spectrum offers both aggregate and individual icons for representing entities in your infrastructure.

More information:

[OneClick Topologies](#) (see page 13)

Aggregate Icons

An IP or physical address does not manage an aggregate icon. However, you can configure aggregate icons to display the device IP address that a container represents. Or, you can configure them to display the subnet address of the devices that the container represents. Aggregate icons primarily act as containers, or placeholders, in a topology view.

An aggregate icon often represents a network group. Some examples of network groups are LAN, LAN_802.x, FDDI, ATM_Network, WA_Link, and Dialup_Link. The following image is an example of an aggregate icon:




CA Spectrum offers many types of aggregate icons. The appearance of an aggregate icon is always based on the entity it represents in your network.

Individual Icons

Individual icons are typically associated with an IP address or a physical address. CA Spectrum can communicate directly with the devices the individual icons represent, so long as the entities they represent are SNMP and ICMP entities.

Individual icons often represent network devices. Some examples of individual icons are those icons that represent a router, switch, or host. The following image is an example of an individual icon:



CA Spectrum offers many types of individual icons. The appearance of an individual icon is always based on the entity it represents in your network.

Icon Types by Theme

You model devices in your network through Discovery or manual modeling. During modeling, CA Spectrum automatically determines the functionality of each device and selects the appropriate icon shape and symbol for that device.

Icons come in various shapes and sizes. Icon symbols vary by the model class that is represented and by the topology in which the icon is located.

Icon types include:

- VNM icon
- Network group icon
- Device icon
- Off-page reference icon
- Segment icon
- Live pipes (links)

VNM

The Virtual Network Machine (VNM) icon typically appears in the top-level view above the network group icons. The background color of the VNM icon changes to indicate the current condition of the SpectroSERVER. For example, the VNM icon turns red when SpectroSERVER disk space reaches 90 percent capacity.



Network Group Icon

The network group icons represent network groupings, for example, cable groups, LANs, and IP Class A, B, C networks. The following icon shows an example of a LAN network group icon:



Device Icon

A device icon represents an individual device. The device icon color changes to indicate the current condition of the modeled device. For example, the device icon changes to red whenever CA Spectrum detects a serious condition requiring attention. The following icon shows an example of a device icon:



Off-Page Reference Icon

The off-page reference icon is a special purpose topology icon. The off-page reference icon represents a device that is directly connected to a device in the current view but which is modeled in another layer. The following icon shows an example of an off-page reference icon:




Note: You can globally suppress off-page references in topology views. For more information about suppressing off-page references, see the *Operator Guide*.

More information:

[Resolve Unresolved Connections](#) (see page 83)

Segment Icon

Segment icons represent conceptual elements of a network. Examples of segment icons can include a coax segment, a wa_segment, a fanout, an unplaced icon, and a pingable. The following icon shows an example of a segment icon:



Live Pipes (Links)

Live pipes represent the connection status between network devices. The links change color to indicate the current condition. A gold pipe represents a resolved connection or indicates that live pipes have not been enabled for the connection. A silver pipe represents an unresolved connection.

Live pipes are not enabled by default. To monitor the connection status between devices, enable a live pipe. The following image shows an example of a live pipe which has been enabled:



More information:

[Connections \(Pipes\) Between Modeled Devices](#) (see page 79)

[Enable or Disable Live Pipes on Individual Links](#) (see page 225)

[Enable or Disable Live Pipes System-Wide](#) (see page 225)


[Enable or Disable a Live Link](#) (see page 84)

Icon Color and Condition


All icons change color to indicate the condition of the device or devices they represent. For instance, a device icon changes color when an alarm condition for that device occurs. A rollup triangle on a device icon or container icon changes color when an alarm condition occurs on one or more of its components. The components include devices or interfaces.

Rollup Condition Colors

The rollup triangle that is associated with the network container icon indicates that one or more components of the container has an alarm condition. In this example, a device within the network container has a critical alarm.



The rollup triangle that is associated with the device icon indicates that a component of that device has an alarm condition. In this example, an interface on a Cisco router has a minor alarm.



Icon Condition Colors



Logical links (or pipes) change color to indicate the condition of the connection:

- Disabled or maintenance conditions = brown
- Good conditions = green
- Initial conditions = blue
- Suppressed or unknown conditions = silver
- Poor conditions = red

Provision Access to Modeled Elements

As an administrator you can secure access to models by applying a security string. A security string establishes permission to various modeled elements in a OneClick topology view such as a modeled device.

After a security string is applied to a modeled device, all subcomponent models (or views) of that device inherit the security string. The security string field for implementing this model security appears in the Component Detail panel, as shown in the following example:

The screenshot shows the 'SPECTRUM Modeling Information' panel with the following fields:

Field	Value	Action
Community String		
Poll Interval (sec)	60	set
Polling	On	set
DCM Timeout (ms)		
DCM Retry Count		
Is a Proxy Model		
Disable Trap-Based Events		
Telnet Port	23	set
SSH Port	22	set
Agent Port		
Security String	Boston	
Landscape	telwin (0x1000000)	
Creation Time	Nov 16, 2007 4:39:31 PM EST	
Device Type		
Model Type Name	VNM	
Model Class	Application	set
Lock Model Class		
System Object ID		

As shown, the security string 'Boston' prevents any user that does not have an Access Group of 'Boston' from accessing this modeled element.

Any user with an Access Group of 'Admin' overrides model security; such users can access all model elements regardless of the security string implemented.

Note: For more information about creating or renaming Access Groups that are associated with individual users or user groups, see the *Administrator Guide*.

Chapter 2: Discovering and Modeling Your Network

This section contains the following topics:

[Methods for Modeling Network Entities](#) (see page 27)
[Discovery](#) (see page 28)
[Discovery Console](#) (see page 30)
[Open the Discovery Console](#) (see page 40)
[Define a Discovery Configuration](#) (see page 41)
[Activate a Discovery Session](#) (see page 56)
[Activate a Modeling Session](#) (see page 57)
[Run a Network Services Discovery](#) (see page 58)
[Create Discovery Configuration Folders](#) (see page 59)
[Reorganize Discovery Configurations](#) (see page 59)
[Rename Discovery Configurations or Folders](#) (see page 60)
[VLAN Discovery](#) (see page 60)
[View, Filter, and Export Results Lists](#) (see page 60)
[Import Discovery Configurations](#) (see page 64)
[Export Discovery Configurations](#) (see page 64)
[After Discovering and Modeling](#) (see page 65)
[VNM AutoDiscovery Control Settings](#) (see page 65)

Methods for Modeling Network Entities

As an administrator, you can define models to represent entities in your IT infrastructure. You define models by manually modeling them or by having Discovery create them for you.

Note: Modeling by Discovery requires less time and less effort than modeling manually.

When you create models in CA Spectrum, consider following these steps:

1. Plan.
2. Use Discovery to model entities.
3. Define or edit models, as needed, using the manual modeling feature.

Create a plan about the network entities you want to model by gathering all required network device information.

When using the automated Discovery process, you must have the following information about your devices:

- IP address ranges of your devices
- Router addresses (optional)
- SNMPv1, v2c SNMP community strings, or SNMPv3 security credentials

When manually modeling your network, you must have details about all the network devices you plan to model, including:

- Device SNMPv1, v2c SNMP community strings, or SNMPv3 security credentials
- Type of network (switched, routed, flat)
- Network masks
- Network technology (FDDI, Ethernet, WAN, and so on)
- List of all devices
- IP addresses for all addressable devices
- Physical and logical network diagrams

Note: All four core topologies in OneClick (Universe, Global Collections, World, and TopOrg) support manual modeling operations. However, we recommend that you always model new devices in the Universe topology and reuse these modeled devices to create other topology views.

Discovery

Discovery is a utility that you can run to find devices in your network and to model them automatically in the Universe topology. Discovery uses a set of configuration parameters that you can modify to determine which network entities to discover and model. You can reuse any set of previously saved configuration parameters and you can also rename, duplicate, or delete configurations.

A configuration determines the focus and scope of a Discovery or modeling session. You define the configuration by selecting parameters on the Discovery Console Configuration tab. After you create a configuration, you can choose when to activate it:

- You can activate the configuration immediately.
- You can schedule the activation, including scheduling it to recur.
- You can save the configuration and activate it later.

Depending on your user privileges, you can use the automated Discovery and modeling features together or you can use them separately. For example, here are some ways you can use Discovery:

- **To perform network inventories:** With read/write privileges to the Discovery parameters, you can use Discovery to identify assets within your network. And, you can export, as needed, the results describing those assets to a desired file format for further review and distribution.
- **To model network entities that you want to manage:** With read/write privileges to both Discovery and modeling parameters, you can use Discovery to:
 1. Determine which elements in your network you want to model.
 2. Identify how you would like CA Spectrum to model these elements for you automatically.

Specifying modeling parameters with the Discovery parameters lets you easily create accurate software models of your infrastructure with less time and effort than modeling manually.

More information:

[Configuration Tab](#) (see page 30)

Separate Configurations

Creating separate Discovery and modeling configurations offers you more flexibility for customizing the Discovery and modeling process. By providing separate configurations, you can:

- Discover limited portions of your network by performing several smaller Discovery operations instead of performing one large Discovery operation.
- Model the results of a Discovery operation using different modeling options.
- Filter and export the results of a Discovery session in different ways.
- Filter and export the results of a modeling session in different ways.

More information:

[Filter, Sort, Export, Search, and Model Discovery Results](#) (see page 35)

[Define Modeling Options](#) (see page 53)

Discovery Console

The Discovery Console consists of two panels: the Navigation panel on the left, and the Contents panel on the right.

Navigation Panel

The Discovery Navigation panel contains the Landscape drop-down list, a toolbar, and a list of configurations and folders available on the selected landscape. From the toolbar you can create, copy, delete, import, or export configurations and create new folders for the configurations. In the Name column, you can select a configuration to open it and view its details in the Contents panel.

Contents Panel

The Discovery Console groups the parameters you use to define Discovery and modeling configurations into four tabs in the Discovery Contents panel:

- [Configuration Tab](#) (see page 30)
- [Discovery Tab](#) (see page 33)
- [Modeling Tab](#) (see page 36)
- [History Tab](#) (see page 38)

Configuration Tab

The Discovery Configuration tab lists all the required and optional parameters you can set to create a configuration.

This tab contains the following settings:

Seed Routers

Specifies the IP addresses or host names of your network seed routers that act as an initial communication point for discovering the network topology. For a host name, CA Spectrum attempts to resolve the host name to an IP address when you start Discovery. If CA Spectrum cannot resolve the host name to an IP address, an error message occurs. The error message displays in the Discovery status panel in the lower section of the Discovery tab. If you start Discovery in the context of a network element, the device IP address is populated in the Seed Router field. You can add more seed router IP addresses as needed.

IP/Host Name Boundary List

Specifies IP ranges, IP addresses, host names, or any combination of this information that CA Spectrum can use to define the boundaries for the configuration. For a host name, CA Spectrum attempts to resolve the host name to an IP address when you start Discovery. If CA Spectrum cannot resolve the host name to an IP address, an error message occurs. The error message displays in the Discovery status panel in the lower section of the Discovery tab. You can start Discovery in the context of a container in the Topology tab or the Explorer tab. In this situation, CA Spectrum populates the IP/Host Name Boundary List with an address range. The address range is determined using the IP address and network mask of the selected device. You can specify more IP addresses, IP address ranges, or host names.

Note: The IP/Host Name Boundary List also accepts single IPv6 addresses; however, IPv6 ranges are not supported.

Import

Imports IP Addresses or host names to populate the IP/Host Name Boundary List. Using a text file, you can import the following information into a configuration:

- One or more single IPv4 addresses
- IPv4 ranges
- Single IPv6 addresses
- One or more host names

You can import IP addresses and host names in a one-time, static manner. Or, you can import the addresses each time the configuration is activated, or dynamically.

SNMP Information

Specifies SNMP community strings and profiles for SNMPv1, SNMPv2c, and SNMPv3 communication.

Modeling Options

Specifies whether to perform a Discover only operation or a combined Discover and model operation. To review and accept the modeling defaults or to edit them, click the Modeling Options button.

Advanced Options

Contains the Advanced Options button which opens the Advanced Options dialog. In the Advanced Options dialog you can review, accept, or redefine the following options:

SNMP Ports

Specifies the default SNMP port and any additional ports. This feature is most often used for managed node environments that use port numbers other than the default port number of 161.

IP Exclusion List

Create, delete, modify, or import an IP exclusion list. This list instructs the Discovery session to exclude the devices in a defined IP address range.

Discovery Options

Specifies whether the Discovery process uses ICMP and Route Tables. For the Route Tables option, you can set a Throttle level to control how often the server sends SNMP requests.

Auto Export

Specifies whether to export the Discovery session results automatically and the preferred format for exporting them: comma-delimited, tab-delimited, or web page.

Scheduling Options

Specifies whether to activate certain configurations regularly using a schedule.

Note: In a DSS environment that spans multiple time zones, the local time of each SpectroSERVER is used for scheduling. For more information about OneClick schedules, see the *Operator Guide*.

Save options as default

Specifies whether to save the current configuration settings as the default configuration. For example, by default the 'Discover only' option is enabled on the Configuration tab. To change the default setting to the 'Discover and automatically model to CA Spectrum' option, select that option. Then, select the 'Save options as default' check box, and select File, Save.

Discover

Activates the Discovery session as it is defined in the Configuration tab.

More information:

[Define Modeling Options](#) (see page 53)

[Discovery Console](#) (see page 30)

[How to Set Discovery Configuration Parameters](#) (see page 43)

[Discovery](#) (see page 28)

Discovery Tab

The Discovery tab displays the results and status of the most recent Discovery session for the configuration that is selected in the Navigation panel. The discovered devices appear at the top of the Discovery tab. In the lower section of the Discovery tab, the status and error messages that are generated display in the Discovery Status panel. All users with Discovery privileges can access this tab. The Discovery tab initially appears disabled for new Discovery configurations and becomes enabled after an initial Discovery session is generated from the Configuration tab.

The Discovery results table includes the following columns by default: Discovered IP, System Name, Device Type, and System Description. To display the Table Preferences dialog, right-click one of the column headers. From this dialog, you can select more columns to display. The Model State column tells you whether the device is modeled in OneClick. This information is useful when identifying devices that are discovered on your network that require modeling in OneClick.

Contents: New Configuration 1 [set](#)

Configuration | **Discovery** | Modeling | History

Filter: Displaying 60 of 60

Discovered IP	System Name	Device Type	System Description
10.253.8.146	Device-2.ca	Router	Cisco IOS Software, RPM Software (RPM-JK9O35-M), Versi
10.253.9.2	rs8022	Router	RS 8000 - Riverstone Networks, Inc. Firmware Version: 8.0
10.253.9.4	10.253.32.1	Router	RS 2000 - Riverstone Networks, Inc. Firmware Version: 8.0
10.253.27.2	xp801	Router	XP-8600 - Enterasys Networks Firmware Version: E9.1.9.4
10.253.48.1	R538	Router	RS 38000 - Riverstone Networks, Inc. Firmware Version: 9.
10.253.48.3	rs8000	Router	RS 8000 - Riverstone Networks, Inc. Firmware Version: 8.0
10.253.158.4		Pingable	
10.253.179.233	Rtr160	Bridge	Software (AbC-I-L)Version 13.0.(5a)Compiled by sclause
10.253.179.235	Rtr730T_248	Bridge	IO SoftwareIOS Software Version 11.3(8)T1Compiled by ps
10.253.180.55	FastIron-1.55	Bridge	Foundry Networks, Inc. FastIron Workgroup Switch, IronW
10.253.180.75	FastIron-1.75	Bridge	Foundry Networks, Inc. FastIron Workgroup Switch, IronW
10.253.190.1	ciscoRSM-9.ca	Router	Cisco Internetwork Operating System Software IOS (tm) C.

Discovery Status:

Starting Range Discovery: Fri Oct 12 10:29:37 EDT 2007
 Discovering 76 single IP addresses - Fri Oct 12 10:29:37 EDT 2007

No devices found at the following IP addresses:

- 10.253.2.21
- 10.253.31.130
- 172.19.2.18
- 172.19.6.5
- 172.19.8.60
- 172.19.11.2

Search: ☒ Highlight All ☒ Ignore Case

More information:

[Discovery Console](#) (see page 30)

Filter, Sort, Export, Search, and Model Discovery Results

The Discovery tab provides the following options to help you review, filter, export, and model Discovery results:

- **Filter:** The Filter text box lets you quickly filter the devices in the results list. For example, to develop, model, and export a list of Cisco devices from your results list, complete the following procedure:
 1. Type **Cisco** in the Filter text box.
This filters the results list by Cisco devices.
 2. Click Model.
This models only the Cisco devices.
 3. Click Export.
This exports the Cisco device results list.
- **Advanced Filter:** To apply more filter criteria, click the Advanced Filter button and creating one or more expressions. These expressions let you set more filters on the Discovery results list.
- **Exclude:** You can exclude one or multiple entries in the Discovery results list by right-clicking the entries that you want to exclude and selecting Exclude. You can also exclude these devices from the Discovery configuration. Select 'Save options as default' in the Configuration tab. Exclude one or more devices, then save the configuration. Discovery excludes those devices when it runs.
- **Export:** Click Export to open the Export table data to file dialog. In this dialog, you can specify a file format and location to export the Discovery results.
- **Status Search:** Enter character strings that you want to search for in the Discovery or the Model Status sections into the Search text box. To see all search matches in the Status panel, select the 'Highlight All' check box. To make the search case-insensitive, select the 'Ignore Case' check box. Use Next and Previous to navigate through the search matches in the Discovery Status panel.
- **Model:** Click to open the Modeling Configuration dialog. In this dialog, you can review or modify the default modeling options provided. When you click OK, CA Spectrum models only the devices appearing in the Discovery results list.

More information:

[Export a Results List](#) (see page 61)

[Define Modeling Options](#) (see page 53)

[Filter Results Using Advanced Filter](#) (see page 62)

[Separate Configurations](#) (see page 29)

[Modeling Tab](#) (see page 36)

Modeling Tab

The Modeling tab displays the results and status of the last modeling session. The top portion of the tab shows the modeled devices. The Modeling Status section at the bottom displays the status and error messages that are generated during the last modeling session.

Contents: New Configuration 1 [set](#)

Configuration | **Discovery** | **Modeling** | History

Filter: Displaying 60 of 60

Condition	Name	Primary Contact IP	Manufacturer	Model Class	Model Status
Normal	Rtr1500_16	10.253.179.233	Cisco System...	Switch-Router	Active
Normal	ciscorpm-9...	10.253.190.1	Cisco System...	Switch-Router	Active
Normal	10.253.32.1	10.253.9.4	Riverstone N...	Switch-Router	Active
Normal	CIAGENT	172.19.10.027	Compaq	Workstation...	Active
Normal	Rtr7301IPT...	10.253.179.235	Cisco System...	Switch-Router	Active
Normal	cat5000-sup...	172.19.94.82	Cisco System...	Switch	Active
Normal	ciscorpm-9...	10.253.8.145	Cisco System...	Switch-Router	Active
Normal	rs8000-18.3	10.253.49.3	Riverstone N...	Switch-Router	Active
Normal	FastIron-18...	10.253.180.75	Foundry Net...	Switch	Active
Normal	cisco7204-9...	172.19.95.5	Cisco System...	Switch-Router	Active
Normal	bn01-sun	172.19.246.98	net-smip	Workstation...	Active
Normal	FastIron-18...	10.253.180.55	Foundry Net...	Switch	Active
Minor	HPAGENT	172.19.246.104	Microsoft	Workstation...	Active

Modeling Status:

Starting Modeling Process: Fri Oct 12 10:30:22 EDT 2007
 Preparing SPECTRUM database for new models
 60 manageable entities destined for management in the SpectroSERVER.
 Model of type Rtr_Cisco at IP 10.253.179.233 is active.
 Identified existing model of type Rtr_Cisco at IP 10.253.179.233
 Mapping Router for model at IP: 10.253.179.233
 Created new model of type Rtr_Cisco at IP 10.253.190.1
 Model of type RstonsSwRtr at IP 10.253.9.4 is active.
 Identified existing model of type RstonsSwRtr at IP 10.253.9.4
 Mapping Router for model at IP: 10.253.9.4
 Model of type Host_Compac at IP 172.19.10.027 is active.
 Identified existing model of type Host_Compac at IP 172.19.10.027
 Router at 10.253.179.233 has been mapped.

Search: ☒ Highlight All ☒ Ignore Case

Model Creation | Activation/Layer 3 | Layer 2 Mapping | Network Services | Auto Placement

All users with modeling privileges can view the Modeling tab. This tab initially appears disabled for new Discovery configurations and becomes enabled only after an initial modeling session is activated from the Discovery tab.

The Modeling tab provides the same options that the Discovery tab does to help you review, filter, export, and search modeling results.

Modeling Session Status Bar

The Modeling Status section displays a status bar at the bottom. The status activates immediately after clicking the Model button on the Discovery tab. This status bar divides the modeling process into four operation phases:

Model Creation

Phase 1—The label for this phase turns green while the SpectroSERVER processes the data to model.

Activation/Layer 3

Phase 2—The label for this phase turns green while the SpectroSERVER maps Layer 3 devices.

Layer 2 Mapping

Phase 3—The label for this phase turns green while the SpectroSERVER waits for model activation and maps Layer 2 devices.

Network Services

Phase 4—The label for this phase turns green while the SpectroSERVER processes the running status for each network service.

Autoplacement

Phase 5—The label for this phase turns green while the SpectroSERVER places the models appearing in the modeling result list.

On the Discovery Console, the Status box at the bottom displays the status and error messages that are related to each of these phases.

More information:

[Discovery Console](#) (see page 30)

[Filter, Sort, Export, Search, and Model Discovery Results](#) (see page 35)

History Tab

The History tab displays information about the configuration that is selected in the Discovery Navigation tab. The following image shows an example of the History tab.

Contents: New Configuration 1 [set](#)

Configuration | Discovery | Modeling | History

Discovery Time	New De...	Lost Devices	Last Time of Discovery Witho...
Oct 1, 2007 9:46:59 AM EDT			
Oct 12, 2007 10:30:16 AM EDT View Changes	3	1	

Discovery Results | Discovery Status | Modeling Status

Filter: Displaying 58 of 58


IP Address	System Name	Device Type	System Description
10.253.158.4		Pingable	
10.42.246.20		Pingable	
172.19.30.0		Pingable	
172.19.4.0		Pingable	
172.19.5.0		Pingable	
172.19.55.0		Pingable	
172.19.57.0		Pingable	
172.19.58.0		Pingable	
172.19.59.0		Pingable	
172.19.6.0		Pingable	
172.19.7.31		Pingable	
172.19.64.51		Bridge	Cisco Systems W5-C5000Cisco Catalyst Operating System
172.19.59.100		Bridge	Ethernet Switch 425-24T HW:06 FW:3.5.0.2 SW:v3.
10.42.94.51	"Cat6505-94.51"	Bridge	Cisco Systems W5-C6506Cisco Catalyst Operating System
10.253.9.4	10.253.32.1	Router	RS 2000 - Riverstone Networks, Inc. Firmware Version: 8.
10.82.246.98	hsun.ca.com	Host	SunOS hsun.ca.com Generic

Discovery Time

Displays the time and date that a Discovery session occurred.

View Changes

Opens the Discovery Result Differences dialog. The dialog lists all devices that were either found or lost during the selected Discovery session when compared to the previous Discovery session. You can filter the information and can export it to a file.

**New Devices**

Displays the number of new devices that are found since the previous Discovery session for the selected configuration.

Lost Devices

Displays the number of devices that are lost between a previous Discovery session and the next Discovery session for the selected configuration.

Last Time of Discovery Without Changes

Displays the time and date information for when the selected configuration last ran without any changes. No time and date information is displayed when changes have occurred each time that the configuration has run.

The History tab also displays the Discovery Results, Discovery Status, and the Modeling Status tabs in the lower panel.

More information:

[Discovery Console](#) (see page 30)

Discovery Connection Status

To display the Discovery Connection Status dialog, click the Connection status icon (). In the Discovery Console, this icon is located in the Status bar.

When displayed from the Discovery Console, the Connection Status dialog shows:

- SNMP Service (SpectroSERVER) connection status
- Web server connection status
- Landscape (SpectroSERVER) connection status

Note: For more information about the Connection Status dialog and the Status bar, see the *Operator Guide*.

Open the Discovery Console

When creating a Discovery configuration, you can open the Discovery Console from the context of a selected model. You can also open the Discovery Console without this context.

To open the Discovery Console without specific model context, do *one* of the following steps:

- Click Tools, Utilities, Discovery Console. The Discovery Console opens.
- Click  (Discovery) in the Topology tab toolbar without selecting a device in the Navigation panel or Topology tab. The Discovery Console opens.
- Click  (Discovery) in the List tab toolbar without selecting a device in the Navigation panel or Topology tab. The Discovery Console opens.

Discovery can also be launched with context. Launching in context automatically generates a Discovery configuration that is based on the models that are selected in the OneClick Console. Discovery retrieves IP/subnet mask information from the selected routers, switch-routers, and LAN models and automatically creates a Discovery configuration that is based on this information.

To open the Discovery Console with context, do *one* of the following steps:

- Right-click a selected model in the Universe Navigation panel and select Utilities, Discovery Console.
- Right-click a selected model in the Contents panel, Topology tab. Select Utilities, Discovery Console.

- Select a model in the Topology tab and click  (Discovery) in the toolbar.
- Right-click a selected model in the Contents panel, List tab. Select Utilities, Discovery Console.
- Select a model in the List tab and click  (Discovery) in the toolbar.

The Discovery Console opens. If you select certain types of containers, a router, or a switch-router, the Configuration dialog also opens from the context of the selected model.

Discovery derives the context differently, based on how you create the Discovery configuration, as follows:

- From the Discovery button (Topology or List tab) – Discovery derives the context using your selection in the Contents panel.
- From the Tools menu – Discovery gets its context from your most recent selection in the Navigation panel.

Define a Discovery Configuration

A Discovery configuration determines which network devices to discover. When creating a Discovery configuration, you can open the Discovery Console from the context of a network element or container. You can also open the Discovery Console without this context.

Without context: Follow these steps:

1. [Open the Discovery Console](#) (see page 40).
2. [Specify the Discovery configuration settings](#) (see page 43).

Note: The IP/Host Name Boundary List and SNMP Information sections are mandatory for a successful Discovery configuration.
3. Do *one* of the following steps:
 - Click Discover to activate a Discovery session for the configuration you created.
The results of the Discovery session appear on the Discovery Console, Discovery tab.
 - Click File, Save.
The configuration that you created is saved.

Note: To apply the most recent changes that are made to any Discovery configuration, select the 'Save options as default' check box.

You can create a Discovery configuration with seed router context. As expected, Discovery discovers the selected router model. Plus, depending on other configuration parameters, Discovery discovers the LANs and routers that are connected to the seed router.

A Discovery configuration with container model context finds all of the network devices that reside within that container IP range.

With context: Follow these steps:

1. [Open the Discovery Console](#) (see page 40).
2. Do *one* of the following steps:

- If the Configuration dialog opens, enter a name for the configuration and click OK.

Note: The Configuration dialog opens when Discovery cannot find an existing Discovery configuration for the selected device.

- If the Use Existing dialog opens, do *one* of the following steps:
 - Select an existing configuration that you want to use and click OK.
 - Click Create, enter a name for the new configuration, and click OK.

If no existing configurations include the IP address of the selected device, the Configuration dialog opens. Use the provided name for the new configuration (which is based on the device type), or enter another name.

- If the Configuration dialog does not open, click  (Discovery). Then, do *one* of the following steps:
 - Select an existing configuration that you want to use and click OK.
 - Enter a name for the new configuration and click OK.

The Configuration dialog closes and the Discovery Console is now fully visible. The Seed Router section of the Configuration tab contains an IP address entry for the selected device. If you select 'LAN containers' in the OneClick Console before launching Discovery, the IP/Host Name Boundary List section can contain one or more IP ranges.

3. [Specify the Discovery configuration fields in the Configuration tab](#) (see page 43).

The Discovery configuration is defined.

How to Set Discovery Configuration Parameters

To set the parameters in the Discovery Configuration tab, do the following steps:

- [Populate the Seed Routers List](#) (see page 43)
- [Specify Host Names and IP Addresses](#) (see page 44)
- [Specify SNMP Information](#) (see page 46)
- [Specify Modeling Options](#) (see page 48)
- [Configure Advanced Options](#) (see page 54)

Seed Routers

The Seed Routers section is optional, but recommended for large Discovery operations. Seed routers are a core list of routers that Discovery uses as a starting point when determining the routed subnets. All routers that are discovered within the IP/Host Name Boundary List are treated as seed routers.

If you start Discovery in the context of a selected device, the IP address for the device appears in the Seed Routers field.

Populate the Seed Routers List

You can populate the Seed Routers list to have Discovery determine the routed subnets from this list.

Follow these steps:

1. In the Seed Routers field, enter addresses or host names to build a list of one or more seed routers and click Add.
2. From the Discovery Type drop-down list, choose *one*:

Routers only

Discovers only the routing devices within the IP range or host name.

Routers and only local LANs in IP/Host Name Boundary List

Discovers only the routed subnets within the IP range or host name.

Routers and all local LANs

Discovers all subnets that are routed by the routers that are discovered in the IP range or host name.

3. Depending on how you want CA Spectrum to discover subnets, choose *one* of the following options:
 - Select the 'Scan Subnets Using ICMP/SNMP Sweep' option and then select the maximum subnet size that you want to discover.
 - Select the 'Discover Subnets Using:' option and then select either 'ARP tables,' 'Cisco CDP tables,' or both.
 - Continue setting configuration parameters in the other view sections, as needed.

Seed Routers information has now been added to the Discovery configuration.

IP/Host Name Boundary List

The IP/Host Name Boundary List section is required for all Discovery configurations. The IP/Host Name Boundary List field populates with an IP range when you start Discovery with a container selected from the Topology tab. The IP range is determined using the network address and network address mask for the container. However, you can specify more IP ranges, individual IPs, or one or more host names, as needed.

- For IPv4, the boundary list accepts single IPs, IP ranges, and host names. Wildcards can also be applied.
- For IPv6, the boundary list accepts single addresses and host names. IPv6 ranges are not valid input. Also, wildcards cannot be applied to IPv6 addresses.
- If an IPv6 address is entered into the first IP address field, the second IP address field is automatically disabled. In this case, you cannot enter an IPv6 range.

More information:

[IP Address Considerations](#) (see page 46)

Specify Host Names and IP Addresses

You can specify which host names, IP addresses, or IP address ranges you want CA Spectrum to discover using these three methods:

- Manual entry
- Import statically
- Import dynamically from a specified file location

Manual: Follow these steps:

1. Enter in the IP/Host Name Boundary List section (first text box) any *one* of the following values:

- A host name
- A single address
- The lowest address in the IP range

Note: You can use a wildcard character to input individual IP addresses. For example, entering 10.10.*.1 could discover: 10.10.0.1; 10.10.16.1; 10.10.32.1; and so on.

2. Enter in the second text box the same single host name or address, or the highest address in the IP range. Click Add.

Important! Attempting to process a large range of IP addresses or several sparsely populated subnets can lead to [unwanted results](#) (see page 46).

3. Repeat Step 1 and Step 2 for each host name, IP address, or IP range of addresses that you want Discovery to contact.

You can also import a list of host names, IP addresses, or IP address ranges for CA Spectrum to discover.

Import statically: Follow these steps:

1. Click Import in the Configuration tab.
2. Select Local Host from the 'Import file location' drop-down list.
3. Select the text file containing the host names or IP addresses. Click Open.

OneClick Discovery reads the host names or IP addresses information from the selected file.

Import dynamically: Follow these steps:

1. Click Import in the Configuration tab.
2. Select 'One Click web server host' from the 'Import file location' drop-down list.

3. Enter the path to the text file containing the host names or IP addresses.

The file must be on the OneClick web server host. The file path must use the native format of the OneClick web server operating system (OS):

- If the web server is running a Microsoft Windows OS, the format of the path must be:

`C:\Program Files\Spectrum\IP_Files\core_network_ips`

- If the web server is running a Solaris or Linux OS, the format of the path must be:

`/usr/Spectrum/IP_Files/core_network_ips`

Note: By default, the installation path for the OneClick web server appears in the Import file path field.

4. Click Open.

OneClick reads the file and imports the host names or IP addresses during each configuration activation. The host names or IP addresses in the text file can be updated regularly, and the updates are reflected in each activation. Importing dynamically lets you maintain current host names or IP addresses automatically when activating a configuration on a scheduled basis.

IP Address Considerations

When specifying IP addresses in the IP/Host Name Boundary List fields, consider the following situations, which can lead to unwanted results:

- Attempting to process a large range of IP addresses. Consider the inclusive range when entering IP address boundaries, and be as specific as possible. You get better results by entering multiple smaller and more pertinent ranges than a single large range. For example, do *not* attempt to run a Discovery with a single IP range of 0.0.0.1 to 255.255.255.255.
- Attempting to discover several sparsely populated subnets. This situation takes significant time because of the timeouts and retries for each unused address. Although many threads are involved in this process, a sparsely populated subnet can quickly exhaust all the available threads. This situation causes the discovery process to take a long time. In this case, [seed router discovery](#) (see page 43) can be a better choice.

SNMP Information

The SNMP Information section is mandatory for Discovery configurations. Here you can review, edit, or remove SNMPv1 and v2c SNMP community strings and v3 security parameters that are currently applied to the configuration. You can also add new strings and profiles manually or you can import them from a text file.

You cannot create a profile name using Import. Instead, first create any desired profile names using the Edit SNMP v3 Profiles dialog before importing the text file. When importing a text file, CA Spectrum compares the SNMPv3 profile names to the existing profiles. These existing profiles were created manually using the Edit SNMP v3 Profiles dialog. If a profile name included in the text file being imported does not exist in CA Spectrum, an error message displays, and the import action fails.

More information:

[Manually Model an SNMPv3 Device](#) (see page 240)

Specify SNMP Information

The SNMP Information section is mandatory for all Discovery configurations. You can manually specify an ordered list of SNMP community strings and profiles for SNMPv1, v2c, and v3. Or, you can import a list of strings. By default, Discovery uses 'public' if no other SNMP community strings are specified.

Note: [For SNMPv3 communication](#) (see page 240), use profiles.

Manual: Follow these steps:

1. Select either the SNMP v1 option or the SNMP v2c option in the SNMP Information section.
2. Type the SNMP community string name for the devices you want discovered in the SNMP Community String field and click Add.

The SNMP community string is inserted into the available SNMP community strings and profile names list.

Import: Follow these steps:

1. Create and save a text file containing the SNMP community strings that you want to use for SNMP. Be sure to use the [correct syntax](#) (see page 47).
2. Click Import in the SNMP Information section.
3. Select the text file that contains the SNMP community strings you want to import.

Valid SNMP community strings and profile names that are imported are added to the available SNMP community strings and profile names list.

Syntax for Imported SNMP Communities and Profiles

You can use this syntax to create a text file that contains SNMP community strings for importing into Discovery configurations.

This syntax has the following format:

```
<name>,<SNMP_version>
```

The text file must list each SNMP community string and profile on a separate line.

<name>

Defines the SNMP community string.

<SNMP_version>

Defines the applicable SNMP version, either 1, 2, or 3.

Note: The version number for SNMPv1 community strings is optional.

Examples: SNMPv1

```
public  
public,v1
```

Example: SNMPv2

```
public,v2
```

Example: SNMPv3

```
public,v3
```

Modeling Options

Modeling configuration settings determine how Discovery models discovered devices. By default, OneClick provides modeling configuration parameters that you can use and modify. Access these settings by clicking the Modeling Options button on the Discovery Configuration tab. Or click the Model button on the Discovery tab.

This Modeling Configuration dialog contains the following settings for configuring discovered models:

Destination Container

Specifies the topology view container where Discovery places discovered device models. You can select a container, such as a LAN container or the New Devices container.

Default: Universe container

Modeling Layout

Specifies how Discovery places and arranges models in the Universe topology view.

Placement

Specifies where models appear in the topology:

- **Flat:** Discovery places all devices, including Layer 1 and Layer 2 in the Destination Container; no LAN containers are created.
- **Hierarchical:** (Default) Discovery places all Layer 3 devices, LAN containers, and Wide Area Links in the Destination Container. Layer 1 and Layer 2 devices are placed in the proper LAN container (based on IP address) under the Destination Container. If Discovery cannot find the appropriate LAN container for Layer 1 or Layer 2 devices, these devices are placed in the Destination Container.

Arrangement

Specifies how models are arranged in the topology:

- Grid
- Radial (Default)
- Tree

Modeling Options

Specifies how Discovery models discovered devices:

Create Wide Area Link Models

Determines whether Discovery creates a WA_Link model between the wide area linked interfaces of two routers. When this option is disabled, Discovery directly connects the linked interfaces.

Default: Enabled

Create LANs (IP Subnets)

Determines whether Discovery uses a LAN container when representing an IP Subnet. Discovery creates the LAN container during the Layer 3 mapping process for any router interface that routes to a local LAN.

Default: Enabled

Remove Empty LANs

Determines whether Discovery destroys any empty LAN containers that the Create LANs (IP Subnets) option created.

Default: Disabled

Create “802.3” (Fanout)

Determines whether Discovery models an 802.3 Fanout segment when CA Spectrum cannot make an accurate connection among three or more interfaces. This model represents the ambiguous connections among these interfaces. However, CA Spectrum uses network traffic data (IfInOctet and ifOutOctet statistics) when the Traffic Resolution protocol option is enabled. This protocol determines connections between interfaces and, in many cases, eliminates the need to model a Fanout.

Default: Disabled

Note: If you have 50 or more connections to a single Fanout model, consider changing this model to a Shared Media Link. The Shared Media Links must be modeled manually. These models can provide more control over fault management behavior when multiple connections are monitored. Unlike a Fanout model, Shared Media Links provide configurable thresholds for handling downstream connections that report problems. For example, a Fanout model reports a problem only when *all* downstream connections are down. However, a Shared Media Link can report the problem sooner, as when 60 percent of the downstream connections are down.

Create Physical Addresses

Determines whether to create a physical address model for a MAC address that is heard by a switch but not associated with any modeled device. The layer 2 mapper attempts to find a connection for each address found. If a connection is found, a Fanout is created and the physical address is associated to it through Connects_To. If no connection is found, the model is placed in Lost and Found. This option is not recommended.

Default: Disabled

New Devices in Maintenance Mode

Determines whether CA Spectrum places the newly discovered devices directly into maintenance mode.

Default: Disabled

Activation Timeout (in minutes)

Determines the number of minutes that Discovery waits for new models to activate before mapping their connectivity. When the timeout expires without any new devices activating, connectivity is established to the extent possible. Connectivity occurs regardless of whether all connections to discovered devices have activated. The minimum activation time is 5 minutes and the maximum time is 15 minutes.

Protocol Options

Lets you configure options for mapping the connectivity between models. By default, Discovery enables several protocol options that are based on best practices. You can disable the default settings or can enable other protocol settings.

Note: IPv6 MIB data is not used for connection mapping.

IP Address Tables

Determines whether to use IP address tables when mapping. When disabled, this option causes Discovery to disable Layer 3 mapping and to map only the Layer 2 connections. Plus, Discovery automatically disables the following options: IP Route Tables Protocol, Create Wide Area Link Model, Create LANs (IPSubnets), and Remove Empty LANs.

Default: Enabled

IP Route Tables

Determines whether to use the IP Address Table when mapping routers. This option is set to No by default because these tables can be large and can take time to read. When this option is enabled, CA Spectrum cannot map unnumbered IP interfaces (0.0.0.0).

Default: Disabled

Source Address Tables

Determines whether Discovery uses the device Source Address table when mapping layer 2 connectivity.

Default: Enabled

Discovery Protocol Tables

Determines whether Discovery uses the Discovery Protocol tables when mapping device connectivity. Supported discovery protocols include Cisco, Nortel, Cabletron Switch, Extreme, Alcatel, Foundry, and Link Layer.

Default: Enabled

ARP Tables for Pingables

Determines whether Discovery uses the ARP table when determining pingable MAC addresses for the connectivity mapping.

Default: Enabled

Spanning Tree

Determines whether Discovery uses the device Spanning Tree Address table (SAT) when mapping Layer 2 connectivity information about the device.

Default: Enabled

Traffic Resolution

Determines whether Discovery uses network traffic data when determining connections between interfaces. In many cases, letting Discovery use the traffic data eliminates the need to model Fanout segments.

Default: Enabled

ATM Protocols

Determines whether the ATM Discovery runs against all ATM switches in the SpectroSERVER database.

Note: For more information, see the *ATM Circuit Manager User Guide*.

Default: Disabled

Network Services Options

Let you specify the network services to run during the modeling process. The supported options include VPN, Enterprise VPN, QoS, Multicast, VPLS, and MPLS Transport. Options are available depending on the components that you have installed.

Filter

Opens the Advanced Filter dialog, where you can exclude selected discovered devices from modeling. You can also click Show Advanced to create a complex filter criterion that includes a combination of AND/OR clauses. The 'Hints' link on the Advanced Filter dialog provides more information.

Auto Export

Specifies whether and how you want to export modeling results automatically. Also lets you select the format for exporting them: comma-delimited, tab-delimited, or web page.

Reset Defaults

Instructs Discovery to use the default modeling settings that are provided with CA Spectrum.

More information:

[Set a Modeling Results List to Export Automatically](#) (see page 61)

[Filter Results Using Advanced Filter](#) (see page 62)

Define Modeling Options

The modeling configuration settings determine how Discovery models the devices it discovers. By default, OneClick provides modeling configuration parameters that you can use or modify. At any time, you can review or change the modeling configuration. Change the configuration using one of these methods:

- Click the Modeling Options button on the Discovery Configuration tab.
- Click the Model button on the Discovery tab.

Note: Before you can define a modeling configuration, [define a Discovery configuration](#) (see page 41).

Follow these steps:

1. In the Discovery Console, with the current Discovery configuration selected in the Navigation panel, do *one* of the following steps:
 - Define a combined Discovery and modeling session by selecting 'Discover and automatically model to CA Spectrum.' To review or modify the modeling configuration, click the Modeling Options button in the Configuration tab.
 - Define a Modeling session after activating a Discovery session by clicking the Model button in the Discovery tab.

The Modeling Configuration dialog opens.

2. In the Modeling Configuration dialog, review or modify any of the fields as needed.
3. How you accessed the Modeling Configuration dialog in Step 1 determines which of these Discovery options you can select:
 - If you clicked Modeling Options on the Discovery Configuration tab to access the Modeling Configuration dialog, click OK. All of your changes are saved, and the Modeling Configuration dialog closes.

Upon activating the next modeling session, Discovery uses the last saved modeling parameters.
 - If you clicked Model on the Discovery tab to access the Modeling Configuration dialog, click OK. The currently specified modeling configuration parameters are saved, modeling session is activated, and the Modeling Configuration dialog closes.

Discovery starts a modeling session that is based on the parameters you have specified.

More information:

[Configuration Tab](#) (see page 30)

[Modeling Options](#) (see page 48)

[Separate Configurations](#) (see page 29)

Advanced Options

Configure Advanced Options if you want to perform any of these Discovery configuration procedures:

- Define the SNMP ports in addition to the default port (161).
- Exclude certain IP addresses from the Discovery.
- Modify the default settings for ICMP, route tables, throttle, time-out, and retries.
- Enable or disable the automatic export of Discovery results.

Follow these steps:

1. Click Advanced Options.
The Advanced Options dialog opens.
2. (Optional) Enter a new port number in the SNMP Ports text box and click Add.
The new SNMP port appears in the list of ports Discovery uses for this configuration.
3. (Optional) Enter an IP address or range to the IP Exclusion List text box.
The IP addresses you enter are excluded from this Discovery.
4. Modify the default Discovery Options settings as needed. These options are as follows:

Use ICMP first, then SNMP

When this option is enabled, Discovery uses ICMP when discovering devices. If ICMP is enabled, Discovery pings the devices in the ranges/subnets first. The devices that responded to ICMP are then queried using SNMP. This option can help reduce the number of SNMP requests, especially when multiple SNMP community strings are being used.

Default: Enabled

Use Route Tables

Use this option only if seed routers are specified in the Discovery configuration. When this option is enabled, Discovery finds neighbor routers and the routed subnets from the IP route tables.

Default: Enabled

Require discovered IP entry in device's IP Address Table

When this option is enabled, Discovery includes only those devices which have the discovered IP address present in their IP address table.

To discover devices that do not have their discovered IP address in the IP address table (e.g. devices using NAT addresses) this option must be disabled. Disabling this option ignores IP address table checking for range discovery.

Default: Enabled

Throttle

Most often this option applies to networks with routing tables containing more than 1,000 entries. If you have networks with routing tables containing over 1,000 entries, you can specify a throttle value (Low, Medium, or High) to stagger the processing workload by having CA Spectrum pause for one second after reading every 50 entries (High), 100 entries (Medium), or 250 entries (Low).

Default: None

Timeout (in seconds)

Specifies the number of seconds that Discovery waits for a device response.

Default: 3

Retries

Specifies the number of attempts that Discovery makes after the first attempt times out before establishing contact.

Default: 3

ICMP Payload Size

Specifies the payload size in bytes. This option is only available if you have selected 'Use ICMP, then SNMP.'

Default: 8

Secure Domain

This option is only available if you have the Secure Domain Manager installed.

Note: For information about the Secure Domain option, see the *Secure Domain Manager User Guide*.

5. Do *one* of the following steps in the Auto Export section:

- To disable Auto Export, select 'Do not export results' from the Auto Export drop-down list.

Auto Export is not enabled for this Discovery and results are not exported automatically.

- To enable Auto Export, select *one* of the following options from the Auto Export drop-down list:

- Export results as CSV (Comma delimited)
- Export results as text (Tab delimited)
- Export results as a web page

Auto Export is enabled for this Discovery. The Discovery results are sent to the location identified on the dialog and in the format that you selected.

6. Click Close.

The Advanced Options dialog closes and your settings are saved to this Discovery configuration.

More information:

[Specify Host Names and IP Addresses](#) (see page 44)

Scheduling Options

The scheduling options are as follows:

- Select an existing schedule to run the configuration (click the Select button in the Scheduling Options section).
- Create a schedule.

If the configuration is scheduled, the schedule name displays next in the Scheduling Options section.

Note: For more information about setting schedules, see the *Operator Guide*.

Activate a Discovery Session

If you have Discovery operations privileges, you can activate a Discovery session by clicking Discover in the Configuration tab. In addition, if you have privileges to modeling operations, you can click Model to activate a Discovery session from the Modeling tab. The following procedure provides instructions for activating a Discovery session for an existing Discovery configuration using the Discovery Console.

Note: If you do not have a Discovery configuration that is ready to activate, Define a Discovery Configuration before activating a Discovery session. We recommend that you also review and make any necessary changes to the Discovery configuration before using this procedure to activate a Discovery session.

Follow these steps:

1. Click Tools, Utilities, Discovery Console.
The Discovery Console opens.
2. Click the name of the Discovery configuration for which you want to activate a Discovery session in the Discovery Navigation panel.
3. In the Discovery Configuration tab, click Discover to do either one of the following tasks:

- To activate a Discovery session and/or combined Discovery and Modeling session.

Discovery activates a Discovery session or a combined Discovery and modeling session that is based on the parameters that are specified in the Configuration tab. The results of the Discovery session appear in the Discovery tab. The results of the modeling session appear in the Modeling tab.

- To rediscover an existing discovered or modeled configuration.

All newly discovered results appear in the results list on the Discovery tab.

Note: Results from this new Discovery session overwrite the results of the previous Discovery session.

Activate a Modeling Session

You can model your Discovery results during the Discovery session, or you can save the configuration and model the results later. You can specify how CA Spectrum models discovered devices that appear in the Discovery tab results list. In the Modeling Configuration dialog, you can accept the default modeling options, or you can change them to meet your requirements.

You can activate a modeling session in the following two ways:

- Select the 'Discover and automatically model to CA Spectrum' option in the Configuration tab before Discovery runs. A modeling session runs automatically with this option selected.
- Click Model in the Discovery tab after a 'Discover only' session has run for that configuration.

Prerequisites to Activating a Modeling Session

- Activate at least one Discovery session of an existing Discovery configuration.
- Review and [make any necessary changes to the modeling configuration](#) (see page 53).
- (Optional) [Exclude certain devices from being modeled](#) (see page 61).
- Ensure that you have sufficient privileges for activating a modeling session.

Follow these steps:

1. In the Discovery Console, click the name of the Discovery configuration for which you want to activate a modeling session in the Navigation panel.
2. Do *one* of the following tasks:
 - In the Configuration tab, select the 'Discover and automatically model to CA Spectrum' option, and click Discover.

Discovery activates a Discovery session and then automatically models the discovered devices appearing in the Results tab of the Discovery tab.
 - In the Discovery tab, click Model to activate a modeling session and model the last set of discovered devices. The last set of discovered devices appears in the Results list of the Discovery tab.

Run a Network Services Discovery

You can discover devices for the following network services: VPN, Enterprise VPN, QoS, Multicast, VPLS, and MPLS Transport. Options are available depending upon the products you have installed.

Note: You can only run one Network Services Discovery at a time.

Follow these steps:

1. Select the models that you want to run a Network Services Discovery for:
 - a. Expand any subnet, folder, and so on, in the Explorer tab or in the List tab, that contains the models you want to select.
 - b. Click the List tab in the Contents panel and select the models by using CTRL+click.
2. Click Tools, Utilities, Network Services Discoveries, and then select the Network Services Discovery that you want to run.

The Discovery executes for the selected models. A pop-up window appears, indicating success or failure of starting the Discovery. The Network Services Discovery fails only when a Discovery is already running for the selected product.

The Discovery process can run for some time. You can check the status of the Discovery process:

- a. Select the product in the Navigation panel.
- b. Click the Information tab in the Contents panel.
- c. Expand the Configuration subview.
- d. Expand the Discovery subview.

Note: For more information about running a specific Network Services Discovery, see the guide for VPN Manager, Enterprise VPN Manager, QoS Manager, Multicast Manager or MPLS Transport Manager. For more information about VPLS, see the *Enterprise VPN Manager User Guide*.

Create Discovery Configuration Folders

Discovery configurations are stored in folders. You can create a folder for a new configuration or a copy of an existing configuration.

Follow these steps:

1. Click Tools, Utilities, Discovery Console.

The Discovery Console opens.

2. Click  (Creates a new folder).

The New Folder dialog opens.

3. Type a name for this Discovery configuration.

The new folder is displayed in the Navigation panel.

Reorganize Discovery Configurations

You can move configurations and folders to different folders using the drag-and-drop method.

Follow these steps:

1. Click Tools, Utilities, Discovery Console.

The Discovery Console opens.

2. Select the configuration or folder you want to move and drag and drop it to the desired location.

The configuration or folder is moved to the desired location.

Rename Discovery Configurations or Folders

You can change the original name of a Discovery configuration after a Discovery session has ended and you can rename Discovery configuration folders.

Follow these steps:

1. Click Tools, Utilities, Discovery Console.
The Discovery Console opens.
2. Right-click the configuration or folder you want to rename and select the Rename node option.
3. Type the new name and click OK.

The configuration or folder you selected displays its new name.

Note: You can also rename configurations by clicking 'set' beside the current configuration name in the Contents panel. However, you cannot rename folders from the Contents panel.

VLAN Discovery

To discover Virtual Local Area Networks (VLANs) successfully, create a root container for each VLAN domain (in a VLAN domain, each VLAN ID is unique). All the devices in a VLAN domain must be placed in a unique root container.

If the devices from different VLAN domains are placed in the same container, VLAN Discovery on that container might not work properly. Also, the VLAN spotlighting might not be able to distinguish the VLANs and devices from the different VLAN domains.

View, Filter, and Export Results Lists

Each time that you activate a Discovery or modeling session, Discovery automatically places the results in the Discovery tab or the Modeling tab. You can use these results to select the devices you want to model or export. For more information, see the following sections:

- [Filter, Sort, Export, Search, and Model Discovery Results](#) (see page 35)
- [Export a Results List](#) (see page 61)
- [Set a Modeling Results List to Export Automatically](#) (see page 61)
- [Filter Results Using Advanced Filter](#) (see page 62)

Export a Results List

You can export a results list by clicking the Export button that appears on the Discovery or Modeling tab. The Export feature accesses the Export table data to file dialog where you can identify:

- Location to save the exported data file
- Name for the exported data file
- Type of file to use to export the data
- Type of file format to save the file

Follow these steps:

1. In the Discovery Console, click the Discovery tab or the Modeling tab.
2. Click the Export button.

The 'Export table data to file' dialog opens.

3. Complete the following information:

Save in

Specifies the location to save the exported data file.

Save as type

Specifies the file type that you want to use when saving the exported data.

File name

Defines the name for the exported data file.

Files of type

Specifies the type of file format to use.

4. Click Save.

The data is exported to the specified location, file name, and file format.

Set a Modeling Results List to Export Automatically

You can automatically export your Discovery or modeling results and status to a web server location.

Follow these steps:

1. In the Discovery Console, select the Discovery configuration for which you want to automate exporting of modeling results.
2. In the Configuration tab, select the 'Discover and automatically model to CA Spectrum' option.

3. Click the Modeling Options button.
The Modeling Configuration dialog opens.
4. In the Auto Export section, do these steps:
 - a. To export modeling results tables, select the 'results table' check box.
 - b. Choose the exported files format from the results tables drop-down list: CSV (Comma delimited), text (Tab delimited), and web page.
 - c. To export status data, select the 'status information (plain text only)' check box.
5. Click OK.

Filter Results Using Advanced Filter

Using the Advanced Filter dialog, you can create filters with compound clauses to exclude certain entries from appearing in the Discovery or modeling results list. If you have privileges to both Discovery and modeling operations, you can access the Advanced Filter dialog before activating a combined Discovery and modeling session. If you have privileges to the Discovery operations only, you can access the Advanced Filter dialog after initiating a Discovery session.

Note: Discovery uses the results list on the Discovery tab to determine which devices to model or export.

Follow these steps:

1. Before you perform a Discovery, do these steps:
 - a. In the Discovery Console, in the Configuration tab, select the Discover and automatically model to CA Spectrum option, and then click Modeling Options.
The Modeling Configuration dialog opens.
 - b. Click the Filter Options button.
The Advanced Filter dialog opens.
 - c. Go to Step 3.
2. After a Discovery session, do these steps:
 - a. In the Discovery Console, click the Discovery tab.
 - b. In the Discovery tab, click the Advanced Filter button.
The Advanced Filter dialog opens.
 - c. Go to Step 3.

3. In the Advanced Filter dialog, complete the fields as follows to create a single expression filter.

Attribute/Ignore Case

Select a device attribute to filter.

Note: If you choose an alphabetic attribute value, you can either clear (ignore) or select (include) the Ignore Case check box.

Comparison Type

Select the type of comparison to make against the value that is specified in the Attribute field.

Attribute Value

Enter or select the desired attribute value to filter.

4. To filter a single expression, click OK.

The Advanced Filter option excludes entities in the results list using the filter parameters that you specified.

Build a compound clause:

5. Click Show Advanced.

The compound expression box and logical operator buttons appear.

6. Click Add to move the single expression (created in Step 2) to the compound expression box.

7. Click *one* of the following logical operator buttons to build a compound expression: New AND; New OR; or AND/OR.

- The compound expression is represented in a tree structure that is grouped by logical operators (AND/OR). Each logical operator in the tree can include any number of attribute criteria nodes and logical operator nodes. For more information, click Hints.
- Alternatively, you can create advanced search expressions using prefix notation.

8. Repeat Step 5 and Step 6 for each compound expression you want to build.

9. Click OK after building the expressions.

The Advanced Filter mechanism excludes the entities in the results list using the compound filter expressions applied.

Import Discovery Configurations

You can import multiple Discovery configurations from your computer, in XML format, into CA Spectrum.

To import a Discovery configuration

1. Click Tools, Utilities, Discovery Console.

The Discovery Console opens.



2. Click (Import).

The Open dialog opens.

3. Browse for the Discovery configuration you want to import from your local computer and click Open.

The Discovery configuration is imported.

Export Discovery Configurations

You can export multiple Discovery configurations, in XML format, to your computer.

Follow these steps:

1. Click Tools, Utilities, Discovery Console.

The Discovery Console opens.

2. Expand the folder where the configuration you want to export is located in the Navigation panel and select the configuration.

3. The configuration information appears in the Contents panel.

Note: You can select multiple configurations. If you select a folder with multiple configurations, all of those configurations are exported. However, the folder hierarchy is not exported.



4. Click (Export).

The Export to file dialog opens.

5. Select the export location on your local computer, enter a configuration name in the File name field, and click Save.

The Export Results from Export dialog opens.

6. Click OK.

The Discovery configuration is exported.

After Discovering and Modeling

CA Spectrum does not support all possible network devices with model types and management modules. To get the CA Spectrum environment you need for managing your network, you can modify the Discovery and Modeling results.

After you successfully discovered and modeled your network, you can examine and enhance the results, for example:

- Modify the device type names to reflect the devices on your network accurately.
Note: For more information about device types, see the *Certification User Guide*.
- Modify the attributes for device models and model types using the Attribute Editor or the Attributes tab.
- Create model types for devices that the CA Spectrum model types and management modules do not directly support.
Note: For more information, see the *Certification User Guide*.
- Import MIBs using the OneClick MIB Tools utility to get updated MIBs for devices, and for features that the provided CA Spectrum MIBs do not support.
Note: For more information about using MIB Tools, see the *Certification User Guide*.
- Modify the names of the existing Discovery configurations.

More information:

[Model Attributes](#) (see page 147)

[Rename Discovery Configurations or Folders](#) (see page 60)

VNM AutoDiscovery Control Settings

AutoDiscovery control settings available on the VNM Information tab affect some of the actions that occur during Discovery and Modeling sessions. If you have a Distributed SpectroSERVER (DSS) environment with multiple SpectroSERVERs, apply all settings changes to *all* of your SpectroSERVERs.

More information:

[AutoDiscovery Control Subview](#) (see page 157)

Access VNM AutoDiscovery Control

To access the VNM AutoDiscovery Control subview, select the VNM in the Explorer tab or in a Universe topology view. Then select the Information tab in the Component Detail panel.

Loopback Interfaces and Discovery

You can set CA Spectrum to use a loopback interface as a primary agent address. Also, you can specify which loopback interface to use when modeling devices.

More information:

[Device Primary Address](#) (see page 129)

Chapter 3: Modeling Your Network Manually

This section contains the following topics:

[When to Model Manually in OneClick](#) (see page 67)

[How to Model Manually in the Universe Topology](#) (see page 68)

[Modeling Manually in a Global Collections Topology](#) (see page 87)

[Modeling Manually in the World Topology View](#) (see page 106)

[Modeling Manually in the TopOrg Topology View](#) (see page 109)

[Using Favorites](#) (see page 111)

[Lost and Found Model Information Subview](#) (see page 112)

When to Model Manually in OneClick

You most often perform manual modeling tasks when you want to represent one or more previously modeled Universe topology devices in other OneClick topologies. For example, OneClick topologies such as Global Collections, World, or TopOrg.

You can also model a network device in CA Spectrum manually after using Discovery in the Universe topology. For example, the Discovery feature is unable to discover new devices in your network that are temporarily offline or blocking management communication. To resolve this situation, you could rediscover these new devices later using Discovery, or you could manually add them to the Universe topology.

In the Universe topology, you can also manually:

- Add devices and annotations to existing models.
- Change device configuration information.
- Improve the readability of models by keeping the layers within the Universe topology simple.

CA Spectrum can model Fanouts automatically, but consider manually changing these Fanouts to a Shared Media Link model when the Fanout has more than 50 connections. The Shared Media Link models use configurable thresholds that can provide more control over fault management behavior.

To keep the layers within the Universe topology simple, consider placing routers near the top and grouping devices logically by IP domains.

Additionally, you can manually model one or more network connections between modeled devices appearing in a Universe or Global Collections view. Finally, manually model all container icons within a World or TopOrg topology view.

How to Model Manually in the Universe Topology

When you manually model in the OneClick Universe topology, consider following this process:

1. [Add containers to the Universe topology views](#) (see page 72).
2. [Add network devices to the Universe topology views](#) (see page 73).
3. Create [connections \(pipes\) between the modeled devices](#) (see page 79):
 - [Create resolved connections](#) (see page 81).
 - [Create partially resolved connections](#) (see page 82).
4. [Cut modeled elements from the Universe topology view](#) (see page 86).
5. [Model manually in the Global Collections topology](#) (see page 87).
6. [Export a topology view](#) (see page 87).

Create Model Dialog

The Create Model dialogs includes the Create Model by Type, Create Model by IP, and Create Model by Host Name dialogs. These dialogs could contain the following settings. The settings that you see vary, depending on the model you are creating and which dialog you are using.

Name

Specifies the unique host name for the device you are modeling.

Note: Model by Host Name supports host names that resolve to either an IPv4 or an IPv6 address.

Network Address

Specifies the IPv4 or IPv6 address for this device, which lets CA Spectrum communicate with the device.

SNMP Community String

Specifies the SNMP community string for this device, which lets CA Spectrum communicate with the device.

Note: You can create a model using Create Model by IP or Create Model by Host Name. If you do not specify a value for SNMP Community String or Agent Port, CA Spectrum uses the predefined SNMP credentials. You configure these SNMP credentials in the VNM model's Information tab in OneClick. Navigate to the Modeling and Protocol Options section of the AutoDiscovery Control subview. If the device cannot be contacted using each of the SNMP credentials but it can be contacted using ICMP, a Pingable model is created.

Serial Number

Specifies the serial number for the device you are modeling.

Security String

Specifies the security for the device. Adding a security string prevents certain users from viewing this model.

Subnet Mask

Specifies the device subnet address that this container represents. When specified, the subnet address label appears whenever a user points their mouse to a container icon in a topology view.

Poll Interval (sec)

Specifies the interval at which this device is polled. To increase or decrease the polling interval, change the value, as needed. By default, CA Spectrum polls modeled devices for status updates every 60 seconds (or for some model types every 300 seconds).

With more time between polling intervals, you use less bandwidth for management traffic, but you receive device status updates less frequently.

With less time between polling intervals, you get more frequent updates of device status, but you use more bandwidth. As such, you may want to use the default polling interval (60 seconds) for critical devices and use 600 seconds for less important devices.

Log Ratio

Defines how many times CA Spectrum polls devices for updates before logging the results.

Default: 10

Creation Author

Specifies the name of the author who is modeling this managed device.

Manufacturer

Specifies the manufacturer name of the managed device you are modeling.

Southbound Gateway-specific

For more information about the Southbound Gateway settings, see the *Southbound Gateway Toolkit Guide*.

Unique ID

The Unique Identifier is a composite of up to six variable data items (1-6). The final unique identified string is composed as follows:

<1>_<2>_<3>_<4>_<5>_<6>

If one of the unique identifier components is not provided, it is not included within the composite unique identifier.

Manager Name

If the name of the third-party application does not apply in the list, choose Default. When this attribute is set on the EventAdmin, all EventModels contained with this EventAdmin inherit this attribute.

Event Model Prefix

This field is prepended to the EventModel name for the Event Models contained by this EventAdmin. This behavior provides consistent naming prefixes for all EventModels associated with a particular EventAdmin. This prefix is useful when sorting or filtering various CA Spectrum applications.

Dialup Link Type

Specifies the functional type of the Dialup_Link. Possible types are Dial Backup Link, Primary on Demand Link, and Bandwidth on Demand Link.

Note: For more information about the Dialup_Link settings, see the *Non-Persistent Connections Manager User Guide*.

Dialup Protocol Type

Specifies the protocol type to use on the Dialup_Link. Possible protocol types include Analog, Switch-56, ISDN, and Frame_Relay.

Activation Grace Period (Min)

Specifies the time (in minutes) that is allowed for the secondary link to become active after a primary link failure. If this grace period expires before the secondary link is active, then a red alarm is generated. Only the Dialup_Link models use this field.

Default: 3 minutes

Deactivation Grace Period (Min)

Specifies the time (in minutes) that is allowed for an active secondary link to deactivate after the failed primary link reactivates. If the secondary link is still active after this grace period expires, then a yellow alarm is generated. Only the Dialup_Link models use this field.

Default: 3 minutes

Active Time Until Yellow (Hours)

Specifies the number of hours a backup link can be active before a yellow alarm is generated.

Active Time Until Orange (Hours)

Specifies the number of hours a backup link can be active before an orange alarm is generated.

Active Time Until Red (Hours)

Specifies the number of hours a backup link can be active before a red alarm is generated.

Device Symbol

Specifies what type of icon to use for this device in the topology view.

DCM Timeout (ms)

Specifies how long the SpectroSERVER waits for a device response.

Default: 3000 milliseconds

DCM Retry Count

Specifies the number of times the SpectroSERVER retries to establish device communication after the DCM timeout value expires.

Default: 2

Agent Port

Specifies the SNMP agent port.

Default: 161

Note: You can create a model using Create Model by IP or Create Model by Host Name. If you do not specify a value for SNMP Community String or Agent Port, CA Spectrum uses the predefined SNMP credentials. You configure these SNMP credentials in the OneClick Information tab for the VNM model. Navigate to the Modeling and Protocol Options section of the AutoDiscovery Control subview. If the device cannot be contacted using each of the SNMP credentials but can be contacted using ICMP, a Pingable model is created.

Secure Domain

Specifies a secure domain for this device. Select the applicable domain from the drop-down list.

SNMP Communications Options

Specifies which SNMP protocol this device supports: SNMP v1, SNMP v2c, or SNMP v3. CA Spectrum uses the protocol that you specify here to discover and map this device.

Profiles

Opens the Edit SNMP v3 Profiles dialog from which you can create profiles for SNMP communication.

Discover Connections

Specifies CA Spectrum Discovery behavior. When this option is selected, CA Spectrum discovers the linked connections (pipes) between this device and other devices.

More information:

[Manually Model an SNMPv3 Device](#) (see page 240)

[Add Containers to Universe Topology Views](#) (see page 72)

[Add a Device Using Create Model by IP Address or Create Model by Host Name](#) (see page 75)

[Add a Device Using Create Model by Model Type](#) (see page 74)

[Edit SNMP v3 Profiles Dialog](#) (see page 239)

[Provision Access to Modeled Elements](#) (see page 26)

Add Containers to Universe Topology Views

You can create a container or can use an existing container to represent the group of devices you want to model. You can create containers at any topology level to help reduce the complexity of your topology views. Containers can effectively help you monitor and manage the health of the devices they represent.

You can manually add a container to a Universe topology view by using the Select Model Type dialog. Some examples of containers you can add include LAN, FDDI, Network, or WAN.

Follow these steps:

1. Select the Universe topology view where you want to add a container.
The selected topology view appears in the Topology tab, Contents panel.
2. Click  (Creates a new model by type) in the Topology tab toolbar.
The Select Model Type dialog opens.
3. Click the Containers tab, select the type of container you want to add, and click OK.

Note: You can use the Filter text box to filter the containers list in the Containers tab. For example, enter **LAN** in the Filter text box to filter the container list to show the LAN container types only.

The Create Model of Type dialog opens.

4. Complete the [fields in the dialog](#) (see page 68).
5. Click OK.

The Create Model of Type dialog closes, and OneClick places the newly created network container in the selected Universe topology view.

More information:

[Editing and Enhancing Topology Views](#) (see page 135)

Add Existing Devices to a Container

To add modeled network devices to a container, double-click the container icon and copy and paste modeled devices into the selected container. These modeled devices can come from other Universe topology views, list views, or the Explorer tab. Or, model new devices in this container by using these topology toolbar functions:

- Create a model by model type.
- Create a model by IP address.
- Create a model by host name.

More information:

[Add Network Devices to Universe Topology Views](#) (see page 73)

Add Network Devices to Universe Topology Views

You can manually add network devices to a Universe topology view using the Topology toolbar 'Create model by' functions. You can add one or more devices to a container using these buttons.

The best practice for manually adding a device to the OneClick environment is to use one of these functions:

- Create a model by IP address.
- Create a model by host name.

More information:

[Add a Device Using Create Model by IP Address or Create Model by Host Name](#) (see page 75)

[Add a Device Using Create Model by Model Type](#) (see page 74)

Add a Device Using Create Model by Model Type

The 'Create model by model type' function is considered an advanced function. This function requires that you have an understanding of how network devices are categorized in the SpectroSERVER modeling catalog.

Follow these steps:

1. Select the Universe topology view where you want the new device to appear.
The selected Universe topology view appears in the Topology tab, Contents panel.
2. If you want to place the new device inside a container, double-click the container icon to display the topology view for that container.



3. Click  (Creates a new model by type) in the Topology tab toolbar.

The Select Model Type dialog opens.

4. Click the All Model Types tab.

A list of model types appears.

5. (Optional) Enter text in the Filter field to filter the list.

As you type characters in the Filter field, only the model type names that contain the same string of characters are shown in the list.

6. Select the model type of the device you are adding and click OK.
7. The Create Model Of Type dialog opens.
8. Complete the fields in the dialog.
9. Click OK.

The Create Model of Type dialog closes and OneClick places the newly created device icon in the selected Universe topology view.

More information:

[Create Model Dialog](#) (see page 68)

Add a Device Using Create Model by IP Address or Create Model by Host Name

You can add a device using the 'Create Model by IP' or 'Create Model by Host Name' function.

Follow these steps:

1. Select the Universe topology view where you want the new device to appear.
The selected Universe topology view appears in the Topology tab, Contents panel.
2. To place the new device inside a network group container, double-click the container icon to display its topology view.

3. Click the Create Model by IP down arrow () and select one of these options in the Topology tab toolbar:

- Create by IP
- Create by Host Name

The Create Model dialog opens.

4. Complete the fields in the dialog.
5. Click OK.

The Create Model dialog closes, the Creating Model dialog indicates that the request is processing. When it closes, the newly created device icon is placed into the selected Universe topology view.

More information:

[Create Model Dialog](#) (see page 68)

Manual Modeling Tips

- To move or enhance the appearance of the recently modeled device icon, click the Edit mode button in the Topology tab toolbar.
- To cut/copy/paste the modeled device icon to another Universe topology view, list view, or the Explorer tab, use one of the following cut/copy/paste functions:
 - The Topology tab toolbar
 - The List tab toolbar
 - The right-click menu in the Explorer tab
- To change configuration parameters of a modeled device, select the modeled device and change the appropriate settings in the Component Detail panel, Information tab. For example, configuration parameters can include the SNMP community string, polling interval, logging interval, and security string.

More information:

[Editing and Enhancing Topology Views](#) (see page 135)

Model By Type Preference

The Model By Type preference lets you specify which model types are available when users manually create models by model type in the Topology tab. When the users click

the 'Creates a new model by type' button () in the Topology tab toolbar, the Select Model Type dialog opens. The Select Model Type dialog contains a tab that is called 'All Model Types'. This tab lists all the model types available for them when they create a model. You can prevent model types from appearing in this list by modifying the Model By Type preference.

Configure Availability of Model Types for Manual Modeling

As an administrator, you can configure the Model By Type preference to exclude or include certain model types. This feature controls which model types that users can see when they manually add models in the Topology tab.

Note: For general information about setting preferences and using the Set Preferences dialog, see the *Operator Guide*. For advanced information about setting user preferences, locking preferences, and administrating preferences, see the *Administrator Guide*.

Follow these steps:

1. Click View, Preferences.

The Set Preferences dialog opens.

2. Expand the Topology Tab folder in the Name column and click Model By Type.

The Set Preferences dialog displays the Include Model Types list on the left and the Exclude Model Types list on the right.

Note: By default, the Include Model Types list includes all of the available model types when you first access the Model By Type preference. You have not yet excluded any model types.

3. (Optional) Enter text in the appropriate Filter field to filter the desired list.

As you type characters in the Filter field, only the model type names that contain the same string of characters are shown in the list.

4. Do one:

- From the Include Model Types list, select the model type that you want to *exclude*. Click the right arrow button to move it to the Exclude Model Types list.

Note: To select several model types at once, press the CTRL key and click each model type.

After you click Apply, the selected model type is moved to the Exclude Model Types list. This model type no longer appears in the Select Model Type dialog.

- From the Exclude Model Types list, select the model type that you want to *include*. Click the left arrow button to move it to the Include Model Types list.

After you click Apply, the selected model type is moved to the Include Model Types list. This model type now appears in the Select Model Type dialog.

5. Click Apply.

The model type changes you made are applied.

6. Click OK.

The Set Preferences dialog closes.

More information:

[Add Network Devices to Universe Topology Views](#) (see page 73)

Polling Interval Changes

You can change the polling interval for any device. To change the interval, enter a new value in the Poll Interval (sec) field in the CA Spectrum Modeling Information subview for the device model. You can also use the Attribute Editor.

Note: Polling intervals also apply to application models, many of which have an initial setting of zero, which in effect disables polling. However, the preferred method for disabling polling for any model is to set the Polling Status attribute to Off.

More information:

[Open Attribute Editor](#) (see page 175)

Disable Polling for a Model or a Model Type

To conserve bandwidth, you can increase the default polling intervals for selected models. Additionally, you may decide that the status of certain devices is not worth any polling bandwidth, even at longer intervals. For example, some network administrators choose to disable modeling endpoint devices such as workstations. The reason is that they do not need to receive the alarms that occur each time that these devices are powered down. Therefore, to model the endpoints but conserve bandwidth from network polling traffic, you can disable polling to these models (or to any models). You disable polling by changing the value of the Polling Status attribute to FALSE.

Polling Status Changes

You can change the polling status for any model by turning on or off polling. To change polling status, use the Polling setting in the CA Spectrum Modeling Information subview for the selected model. You can also use the Attribute Editor.

Note: The Polling Status value takes precedence over the Polling Interval value in terms of enabling/disabling various periodic external requests to a model. Setting the Polling Interval to zero automatically changes the Polling Status to Off. If you reset the Polling Status to On, the CA Spectrum inference handlers could still generate requests, even though the Polling Interval is set to zero. However, to enable normal CA Spectrum polling for fault isolation purposes, the Polling Interval would have to be manually reset to a non-zero value.

Poll Devices That Are Down

When contact with a device has been lost, CA Spectrum uses two methods to continue polling the device for a status change. First, CA Spectrum pings the device once every 60 seconds. Second, CA Spectrum sends SNMP requests every third polling interval by default. For example, if the device polling interval is set to 60 seconds, CA Spectrum polls the down device once every 180 seconds.

To change the default interval at which CA Spectrum polls a down device, insert the following syntax into the <SPECROOT>/SS/.vnmrc file:

```
down_device_poll_interval_multiplier=<user_defined_multiplier>
```

For example, if the <user_defined_multiplier>=2 and the device polling interval is 60 seconds, then CA Spectrum polls the down device once every 120 seconds (2*60=120).

Connections (Pipes) Between Modeled Devices

You can depict the physical connections (pipes) between the modeled devices using the Start and the End Connection options in the Topology right-click menu. In OneClick, you can manually create three types of connections between modeled devices:

- **Resolved connection:** (For a fully resolved connection.) A resolved connection occurs when two devices are connected at the port level. For example, Port-A of device one is connected to Port-B of device two.
- **Partially resolved connection:** A partially resolved connection occurs when only one port is known between two devices. You typically create this type of connection when you know only the port of one modeled device. For example, device-one is connected to Port-A of device-two.

When you manually model a partially resolved connection, CA Spectrum attempts to resolve the port connection of the other device. If CA Spectrum succeeds, it represents the connection as a fully resolved connection. You can later determine whether the connection is fully resolved. To check for a resolved connection, click the link within that view and view the Link Information tab in the Component Detail panel.

- **Unresolved connection:** An unresolved connection occurs when two modeled devices (or containers) are not connected in any way at the port level. For example, container-A is connected to container-B.

Make sure that the Live Pipes attribute for each VNM is set to Enable when these conditions exist:

- You have a DSS (distributed SpectroSERVER) environment.
- Any of the connections or pipes that you are creating span between two or more devices that different SpectroSERVERs manage.

More information:

[Create a Resolved Connection](#) (see page 81)

[Create a Partially Resolved Connection](#) (see page 82)

[Live Pipes Subview](#) (see page 161)

[Live Pipes \(Links\)](#) (see page 24)

Dynamic Link Status: Partial or Fully Resolved Connections

After you create a partial or fully resolved connection between two modeled elements, you can monitor the status of that connection. You monitor the status by enabling the connection as a Live Pipe. The color of live pipes in the Universe topology views indicates status information about the connection. For example, good connection conditions are green, bad connection conditions are red, and disabled live-pipe connections are gold.

A live pipe shows a combined status condition for fully resolved connections (two ports). The connection having the most severe condition determines the color of the pipe. A live pipe can generate an alarm when one or both the links it represents goes down.

Note: Initially after modeling a connection, the color of the connection is gold or silver. Gold appears for resolved or partially resolved connections; silver appears for all unresolved connections.

CA Spectrum monitors the Border Gateway Protocol (BGP) peer session under these circumstances:

- Live pipes are enabled on a connection between two routers, or on a connection between a router and a provider cloud.
- The ports that are involved with the connection are participating in a BGP peer session.

More information:

[Enable or Disable Live Pipes on Individual Links](#) (see page 225)

[Enable or Disable Live Pipes System-Wide](#) (see page 225)

[BGP Peer Session Monitoring](#) (see page 84)

Remove Connections from a Universe Topology View

You can delete a connection between two modeled elements in the Universe topology view by right-clicking the pipe and selecting Delete. When a pipe is deleted, CA Spectrum removes all of its associations. If the pipe represents more than a single port connection, CA Spectrum prompts you to confirm the deletion.

Automatic Recreation of Pipes

CA Spectrum automatically recreates pipes when you copy and paste a set of previously connected modeled icons to another Universe topology view or list view. Also, pipes are automatically recreated when you copy and paste the models in the Explorer tab. If you delete one of the connected modeled icons for a view, the pipe is erased. Later, you could copy that device from the Lost and Found view to the original topology view or list view. In this case, OneClick automatically recreates the connection between the two modeled devices.

Create an Unresolved Connection Between Modeled Elements

When you do not know the port connections between two modeled elements that you want to connect, you can create an unresolved connection. When you create an unresolved connection between two modeled elements in the Universe topology view, the pipe representing the connection is silver. You are prevented from enabling that connection as a Live Pipe. However, after you create an unresolved connection between two modeled elements in a Universe topology view, CA Spectrum automatically attempts to resolve the connection between them. If CA Spectrum succeeds in resolving the connection, the pipe representing the connection is gold and it behaves as a resolved connection. You can then proactively monitor the status of that resolved connection by right-clicking the pipe and selecting Live Pipe.

If CA Spectrum is unable to detect at least one port-level connection between two devices, the pipe in the topology view remains as an unresolved connection (silver). You cannot enable Live Pipe on an unresolved connection.

Follow these steps:

1. In a Universe topology view, right-click any modeled element (device or container) and select Start Connection to designate the starting point of a connection.
2. In a Universe topology view, right-click the modeled element (device or element) and select Connect with *<starting point address>* to designate the endpoint of the connection.

CA Spectrum models an unresolved silver-colored pipe between the two modeled devices that are specified. If the connection between the modeled elements spans two separate views, an off-page reference icon appears in the view.

Create a Resolved Connection

When you know the ports of both modeled devices, you can create a port-to-port resolved connection.

Follow these steps:

1. Designate the starting point of a connection in a Universe topology as follows:
 - a. Select a modeled device (such as a switch or router) that contains port interfaces.
 - b. Click the Interfaces tab in the Component Detail panel.
 - c. Right-click a port row in the Interfaces tab and select Start Connection.

2. Designate the endpoint of the connection in a Universe topology as follows:
 - a. Select a modeled device (such as a switch or router) that contains port interfaces.
 - b. Click the Interfaces tab in the Component Detail panel.
 - c. Right-click a port description in the Interfaces tab and select Connect with *<starting point port address>*.

CA Spectrum creates a resolved (gold-colored) pipe between the two modeled icons. If the modeled devices are in separate views, an off-page reference icon appears in the view.
3. To monitor the link status of this connection, right-click the connection and [select Enable/Disable Live Links](#) (see page 84).

More information:

[Dynamic Link Status: Partial or Fully Resolved Connections](#) (see page 79)
[Enable or Disable a Live Link](#) (see page 84)

Create a Partially Resolved Connection

Sometimes, you know only the device port of one of the two modeled devices that you want to connect. In this case, you can create a partially resolved connection.

Follow these steps:

1. Designate the starting point of a connection by following these steps:
 - a. In a Universe topology view, select a modeled device (such as a switch or router) that contains port interfaces.
 - b. In the Component Detail panel, click the Interfaces tab.
 - c. In the Interfaces tab, right-click a port row and select Start Connection.
2. Designate the endpoint of the connection. In a Universe topology view, right-click any modeled element (device or container) with an unknown port address and select Connect with *<starting point modeled port address>*.

CA Spectrum models a partially resolved (gold) pipe between these two modeled devices. If the connection between the modeled devices spans two separate views, an off-page reference icon appears in the view.
3. (Optional) [Monitor the link status of this connection](#) (see page 84).

Note: OneClick automatically attempts to locate the unknown device port. You can verify whether CA Spectrum locates this device port by clicking the link and viewing the Information tab in the Component Detail panel.

More information:

[Dynamic Link Status: Partial or Fully Resolved Connections](#) (see page 79)

Resolve Unresolved Connections

You can resolve unresolved and partially resolved connections from the Interfaces tab.

Follow these steps:

1. Select the device model for which you want to resolve a connection.
2. Click the Interfaces tab in the Component Detail panel.

A warning appears in the Interfaces tab toolbar specifying how many unresolved connections this device has.
3. Right-click an unused interface and select Resolve Connection To, *<model at other end of connection>*.

The warning in the Interfaces toolbar changes to show the revised number of unresolved connections. For example, if you had only one unresolved connection, the warning now disappears. If you had two unresolved connections, the warning now indicates that you have only one unresolved connection.

More information:

[Off-Page Reference Icon](#) (see page 23)

Lock and Unlock Resolved Connections

In OneClick, you can preserve a resolved connection between two modeled devices by locking that connection. When you lock a connection, Discovery does not delete the connection.

Follow these steps:

1. In a Universe topology view, right-click the resolved connection that you want to lock/unlock and select Lock/Unlock Connection.

The Lock/Unlock Connection dialog opens.

2. Select the connections that you want to lock/unlock and click OK.

The resolved connection is locked/unlocked according to your selection.

Enable or Disable a Live Link

You can monitor the link status of any resolved connection that is depicted in the Universe topology views by enabling Live Links. For live links on a partial or fully resolved connection, you can monitor the port connections at either end of the modeled devices.

The status of a live link displays through colors (red for critical, green for good, and so on). In addition, if a port connection within a live link goes down, you can view alarm information about that connection in the Alarms tab of the Contents panel.

You can enable and disable a live link for a partial or fully resolved connection.

Note: You cannot enable live links on a silver (unresolved) connection.

Follow these steps:

1. Right-click the connection that you want to enable/disable as a live link and select Enable/Disable Live Link.
2. Select the connections that you want to change and click OK.

If you enable a connection, the color changes to green (good condition) or red (bad condition). The color of a disabled connection is gold.

More information:

[Live Pipes \(Links\)](#) (see page 24)

BGP Peer Session Monitoring

Border Gateway Protocol (BGP) peer session monitoring polls the status of the peering session between two BGP devices.

CA Spectrum monitors the BGP port peer session status at the polling interval of the port model's Polling_Interval Attribute value under these conditions:

- The BGP peer session monitoring is enabled.
- The live pipe on a BGP peer session port is turned on.

When a monitored BGP peer session is no longer found in the bgpPeerTable MIB table, the result is:

- A 'BGP peer session removed' event is generated.
- The session is no longer monitored.

Consider the following information about alarms and BGP peer sessions:

- When a WA_Link is connected to the ports in a BGP peer session, a BGP peer session down alarm is generated on one of the port models. The session is not generated on the WA_Link.
- If the loss of a BGP Peering session is the root cause of a downstream outage, a critical BGP alarm is generated. This alarm hides the lost contact alarm of the downstream device.
- If a monitored BGP peer session goes down, a single alarm is generated on the BGP peer session port model. If the monitored BGP peer session is between two directly connected routers, the alarm is asserted on the port model on which the outage is first detected.
- If a BGP peer session port is administratively disabled or operationally down, the BGP peer session alarm on the port model becomes a symptomatic alarm of the link condition alarm.
- When a backwards transition trap is received, the BGP MIB is polled to verify that the peer session is established. If the peer session is not established, an alarm is generated on the peer session port model.

Note: For more information about using MIB Tools, see the *Certification User Guide*.

More information:

[Dynamic Link Status: Partial or Fully Resolved Connections](#) (see page 79)
[BGP Manager Subview](#) (see page 164)

Remove Modeled Elements from the Universe Topology View

You can remove modeled elements from the Universe topology view.

Follow these steps:

1. In the Universe topology view, right-click the modeled element, and select Remove.
The Confirm Removal dialog opens.
2. Click Yes.
The element is removed from the topology view. Models are placed in the Lost and Found.

More information:

[Lost and Found Model Information Subview](#) (see page 112)

Delete Modeled Elements from the Universe Topology View

You can delete modeled elements from the Universe topology view.

Follow these steps:

1. In the Universe topology view, right-click the modeled element, and select Delete.

The Confirm Delete dialog opens.

2. Click Yes.

The model is deleted permanently from the system. When deleting containers, models within the containers are placed in the Lost and Found. If those models exist in other topologies like World View and Global Collection, they continue to exist there. In this case, they are not moved to Lost and Found.

More information:

[Lost and Found Model Information Subview](#) (see page 112)

Cut Modeled Elements from the Topology View or List View

You can cut a modeled element from a topology view, list view, or using the right-click menu in the Explorer tab. In this case, OneClick removes the model from the view and places it into the Lost and Found. If desired, you can remove the modeled element from the Lost and Found view as well.

To cut a modeled element from a view, right-click the modeled element and select Cut.

The cut modeled element moves to the Lost and Found view.

Follow these steps:

1. Select LostFound in the Explorer tab in the OneClick Navigation panel.
2. Click the List tab in the Contents panel.
3. Select the elements that you want to remove in the List tab.
4. Right-click and select Delete.

Enhance Topology Views

You can put the current Topology view in Edit mode. Then, you can enhance the Topology view using the tools that are provided in the Edit mode toolbar.

Export a Topology View

You can export any topology view to a PNG file format.

Follow these steps:

1. Navigate to the topology view that you want to export.
2. Click the Export button in the Topology tab toolbar.

The Save As dialog opens.

3. Specify a name and location for the file and click OK.

The topology view is saved as a PNG file.

Note: For more information about exporting views and table data, see the *Operator Guide*.

Modeling Manually in a Global Collections Topology

If you have the Manage Global Collections privilege, you can create a collection in the Global Collections topology. Create a collection from any modeled elements that were previously modeled in one or more Universe topology views.

When you create a collection within the Global Collections topology, provide a name and an owner for that collection and define its members. The owner field is used to indicate who is responsible for the Global Collection. This field is initially set to the OneClick user who created the collection. Select members for a collection either by specifying search criteria or by using the copy and paste feature.

Dynamic Membership

When you use search criteria to define the members of a Global Collection, the members of that Global Collection are considered *dynamic*. They remain in the Global Collection only as long as they meet the specified search criteria.

CA Spectrum automatically removes a modeled element from the Global Collection that no longer meets the original search criteria, using the method that you choose:

- The next time the Global Collection is manually updated
- Periodically, at the next scheduled interval
- Dynamically, as soon as the model no longer meets the Global Collection search criteria
- Scheduled, according to the assigned schedule

The default period for automatic scheduled collection updates is 24 hours. At any time, you can redefine the dynamic members in a Global Collection by editing the specified search criteria.

Note: Schedules can be assigned only after the initial creation of a Global Collection. Schedules are intended for Global Collections with search criteria that have the capacity to result in degraded product performance.

Static Membership

When you use copy/paste or add functions to define members of a collection, the members of that collection are considered static. The static members always remain in a collection until you decide to remove them manually.

Connections Between Modeled Elements (Members)

The connections between modeled elements in a Global Collection topology view and in a Universe topology view are similar. In both views, connections exhibit the same behavior and functionality. In the Global Collections view, you can create partial and fully resolved connections, or unresolved connections (links, pipes). You can monitor the status of any resolved connection using Live Links.

More Information:

[Connections \(Pipes\) Between Modeled Devices](#) (see page 79)

Updating Modeled Elements in Global Collections

CA Spectrum updates all modeled elements in a Global Collection view, using the method that you choose:

- The next time the Global Collection is manually updated

Note: You must have the Update Global Collection Membership privilege to update a Global Collection manually.
- Periodically, at the next scheduled interval

The default period for automatic scheduled Global Collection updates is 24 hours. At any time, you can redefine the dynamic members in a Global Collection by editing the specified search criteria.
- Dynamically, as soon as the model changes to meet or no longer meet the Global Collection search criteria
- Scheduled, according to the assigned schedule

Note: Scheduled updates can only be applied if you have the Schedule Global Collection Updates privilege. Scheduled updates are intended for Global Collections with search criteria that have the capacity to result in degraded product performance. Scheduled updates are not available during the initial creation of a Global Collection. Without the Schedule Global Collection Updates privilege, you cannot change the Global Collection Update Options when a schedule is applied to a Global Collection.

Generate Reports on Collections

You can generate reports on global collections using the Report Manager module. By using the Report Manager module with the OneClick Global Collections topology, you can at any time generate a single report about any one collection. For more information about running reports, see the *Report Manager User Guide*.

Note: The Report Manager module is not included in the CA Spectrum core product line. This module must be purchased separately.

How to Define and Manage Global Collections

When defining a collection of modeled elements in the Global Collection topology, consider following this process:

1. Create a global collection.
Note: A new global collection remains empty until you define its members.
2. Define dynamic members or, when necessary, define static members.
3. Edit the members in a global collection as needed. For dynamic members in a global collection, you can redefine the search criteria. For static members in a global collection, you can delete, or copy/paste members in a global collection, or can add members to a global collection.
4. Create a Global Collection Hierarchy if you want to organize global collections using folders and subfolders.
5. Delete global collections as needed. The modeled elements within a global collection represent copies of modeled elements from the Universe topology. Therefore, Delete removes only the specified global collection and the copies of the modeled elements that global collection represents.

More information:

[Global Collections Topology](#) (see page 16)

Create an Empty Global Collection

You can create an empty global collection when you are unsure of the type of global collection you want to create. To create an empty global collection, provide only the name of the global collection. Save your empty global collection until you are ready to add static members, dynamic members, or both.

Follow these steps:

1. In the Explorer tab of the Navigation panel, right-click Global Collections and select Create Global Collection.

The Create Global Collection dialog opens.

2. Complete the following fields and click OK:

Name

Specifies the name for the global collection.

Note: If your global collection name matches an existing global collection name, a warning appears. Click Yes to continue naming the global collection, or click No to enter another name for the global collection.

Owner

Specifies the user who is responsible for the global collection.

Description

(Optional) Specifies a description for the global collection.

Security String

(Optional) Specifies a security string expression to prevent certain users from viewing the contents of this global collection.

Note: For information about security string expressions, see the *Administrator Guide*.

The global collection is created and appears in the Navigation panel under Global Collections.

More information:

[Provision Access to Modeled Elements](#) (see page 26)

Create a Global Collection of Dynamic Members

You can create a dynamic collection of dynamic members.

Follow these steps:

1. In the Explorer tab of the Navigation panel, right-click Global Collections and select Create Global Collection.

The Create Global Collection dialog opens.

2. Complete the following fields:

Name

Specifies the name for the global collection.

Note: If your global collection name matches an existing global collection name, a warning appears. Click Yes to continue naming the global collection, or click No to enter another name for the global collection.

Owner

Specifies the user who is responsible for the global collection.

Description

(Optional) Specifies a description for the global collection.

Security String

(Optional) Specifies a security string expression to prevent certain users from viewing the contents of this global collection.

Note: For information about security string expressions, see the *Administrator Guide*.

3. Click Search Options.

The Search Options dialog opens.

4. Complete any of the following fields to create a single search expression:

Attribute

Specifies the attribute of a device to filter. From the drop-down list of commonly used attributes, select the attribute that you want to use. The predefined list might not include the attribute that you want. In this case, click Attribute to specify the model type (device, port, or other) and its associated attribute that you want to find.

Note: If you choose an alphabetic attribute value, you can either clear (ignore) or select (include) the Ignore Case check box.

Comparison Type

Specifies the type of comparison to be made against the value that is specified in the Attribute field. Only the comparison types appropriate to the attribute data type are available.

Ignore Case

Determines whether the comparison is case-sensitive. If you do not select this checkbox, the comparison is case-sensitive. This selection is only enabled when it is appropriate for the data type of the attribute you selected.

Attribute Value

Enter the desired attribute value to search.

Devices Only

Specifies that the search results list includes only devices.

5. (Optional) To use a wildcard character or regular expression in the Attribute Value field, select a valid attribute in the Attribute field. Select Matches Pattern in the Comparison Type field. Then, select one of the following options:

Specify Wildcard Now

Lets you search for a value using a wildcard. The following wildcards are available:

*

Matches *any number* of characters.

For example, 'switc*' returns 'switch' and 'switch-router.'

?

Matches any *single* character.

For example, 'switc?' returns 'switch' but it does not return 'switch-router.'

Both wildcards can be used anywhere and in any combination for a wildcard match.

Note: 'Matches Pattern' is not a valid comparison type for all attributes.

Specify RegExp Now

Specifies that you want to create a search using Perl Compatible Regular Expression (PCRE) matching on attributes of the type 'text string'. Text string searches are available only for Matches Pattern comparison types. PCRE matching helps you to find and group models using specific pattern searches that are more advanced than existing searches or wildcard searches can provide.

Note: By default, all users have the privilege to enter regular expressions. Administrators can disable this privilege on a per-user basis.

6. (Optional) To conduct a compound search clause or a single search clause, do *one* of the following steps:
 - To conduct a search that is based on a single expression, click OK.
 - To conduct a search that is based on a compound clause, complete the following steps to build a compound search clause:
 - a. Click Show Advanced. The compound expression box and logical operator buttons appear.
 - b. Click Add to move the single expression you created in Step 4 to the compound expression box.
 - c. Click one of the following logical operator buttons to build a compound expression: New AND; New OR; or AND/OR.

Note: The compound expression is represented in a tree structure that is grouped by logical operators (AND/OR). Each logical operator in the tree can include any number of attribute criteria nodes and logical operator nodes. For more information, click Hints in the Advanced section.
 - d. (Optional) Click the Add Existing button to [create a global collection from an existing attribute-based search, action-based search, or relation-based search](#) (see page 96).
 - e. Repeat Steps 4 through 5 for each compound search expression you want to build.

Click OK.

The Advanced Search mechanism locates and places a copy of all matching modeled elements (previously defined in the Universe topology) into the global collection.

Note: For more information about searches, see the *Administrator Guide*.

7. (Optional) Select the Real-Time Update check box.

This option disables the update interval. Also, it adds or removes models to and from the global collection when the models meet or no longer meet the search criteria.

8. (Optional) Specify a value in the Run search to update Global Collection membership every <> hours field.

This field determines how often you want OneClick to conduct a search to update the dynamic members that are defined in the global collection.

Note: The option to associate a schedule with a global collection is not available during the initial creation of a collection.

9. Click OK.

The Search Options dialog closes and the Create Global Collection dialog opens.

10. Click Landscapes to identify which landscapes you want Search to include when searching models to populate the global collection.

11. Click OK.

The global collection of dynamic members is created.

More information:

[Provision Access to Modeled Elements](#) (see page 26)

[Edit Dynamic Members in Existing Global Collections](#) (see page 101)

[Global Collection Search Recommendations](#) (see page 94)

Global Collection Search Recommendations

The following information provides search criteria recommendations when defining global collections with dynamic members. The order of the criteria can affect the search performance.

The order of attribute criteria is based on two categories: *storage of information* and *data type*.

Storage of information

Order the attributes from least CPU (quickest access) to most CPU (slowest access), as follows:

- Memory flag (least CPU/quickest access)
- Database flag
- Calculated
- External flag (most CPU/slowest access)

Data type

Order the attributes from quickest comparison to slowest comparison, as follows:

- Integer, counter, enumeration, model type handle (quickest comparison)
- IP address, octet string
- Text string (slowest comparison)

Combining the two categories of criteria, the overall attribute placement for complex searches of AND/OR order from top to bottom is as follows:

1. Memory flag
 - a. Integer, counter, enumeration, model type handle
 - b. IP address, octet string
 - c. Text string
2. Database flag
 - a. Integer, counter, enumeration, model type handle
 - b. IP address, octet string
 - c. Text string
3. Calculated
 - a. Integer, counter, enumeration, model type handle
 - b. IP address, octet string
 - c. Text string
4. External Flag
 - a. Integer, counter, enumeration, model type handle
 - b. IP address, octet string
 - c. Text string

Example

You would like to define a global collection containing dynamic members that are based on the following search criteria (in no particular order):

- ifDesc
- Topology model name string
- Network address
- Model type handle

Using the recommended ordering logic, we recommended the following order:

1. Model type handle (memory flag : model type handle)
2. Network address (memory flag/database flag : IP address)
3. Topology model name string (calculated flag : text string)
4. ifDesc (external flag : text string)

More information:

[Provision Access to Modeled Elements](#) (see page 26)

[Edit Dynamic Members in Existing Global Collections](#) (see page 101)

Create a Global Collection of Dynamic Members from an Existing Search

You can create a global collection of dynamic members from an existing attribute-based search, action-based search, or relation-based search. The existing searches can be found in the Locator tab.

Important! You can create a global collection that is based on an existing search. However, the association between the global collection and the existing search is not maintained. After the global collection has been created, any modifications to the search criteria for the global collection must be made in the global collection itself. Any changes to the existing search (in the Locator tab) on which the global collection was originally based are not propagated to the global collection.

Follow these steps:

1. In the Locator tab, locate the existing search that you want to base your global collection on.

2. Right-click the existing search and select Create Global Collection From.

The Create Global Collection dialog opens.

3. Do the following, as necessary. The type of search you choose determines which options appear:

- Enter a value for the search criteria. Options can include Matches Pattern, Equal To, Contains, and Starts With.
- Select the Ignore Case check box to make the search case-insensitive.
- Click Landscapes to identify which landscapes you want the Search to include when searching models to populate the global collection.
- Click the List button and either enter a list of values to include in the search, or click Import to import a list of values. Click OK, and then click OK again.

The Create Global Collection dialog reopens.

4. Complete the following fields:

Name

Specifies the name for the global collection.

Note: If your global collection name matches an existing global collection name, a warning appears. Click Yes to continue naming the global collection, or click No to enter another name for the global collection.

Owner

Specifies the user who is responsible for the global collection.

Description

(Optional) Specifies a description for the global collection.

Security String

(Optional) Specifies a security string expression to prevent certain users from viewing the contents of this global collection.

Note: For information about security string expressions, see the *Administrator Guide*.

5. (Optional) Do the following steps to add another existing attribute-based, action-based, or relation-based search:

- a. Click Search Options, click Show Advanced, and then click the Add Existing button.

The Add Existing Search dialog opens.

Note: Adding an existing search to your custom search copies the existing search and embeds it, as it is now, into your custom search. If you modify the existing search later, your custom search does not change. Your custom search contains only a *copy* of that existing search, as it was when you first copied it.

- b. Select the existing search that contains the criteria you want to add to the current search and click OK.

The Search dialog opens.

c. Do the following steps:

- Enter a value in the provided field or select a value from the drop-down menu, if the drop-down menu is available.
- Click Landscapes to identify which landscapes you want to include when searching models to populate the global collection.

Click OK.

The criteria that you selected is added to the compound expression.

Note: For more information about searches, see the *Administrator Guide*.

6. Click OK.

The Create Global Collection dialog closes and the global collection is created. The global collection appears in the Navigation panel under the Global Collections folder.

More information:

[Create a Global Collection of Dynamic Members](#) (see page 91)

[Provision Access to Modeled Elements](#) (see page 26)

Create a Global Collection of Static Members

You can create a global collection of static members on the fly from a topology view.

Follow these steps:

1. In a topology view, do one of the following steps to designate the modeled devices that you want to add to a global collection:

- To select a single modeled device, right-click a modeled device in the Navigation panel and select Add To, Global Collection(s).

Note: You can right-click a single modeled device in *any* topology view to add the modeled device to a static global collection.

- To select multiple modeled devices in any topology view, do the following steps:
 - a. Press and hold the SHIFT key and individually select the modeled devices.
 - b. While the SHIFT key is pressed, right-click the last selected modeled element and select Add To, Global Collection(s).

Note: You can also select one or more modeled elements in the List tab and add them to a global collection.

The Select Global Collections dialog opens.

2. Click Create.

The Create Global Collection dialog opens.

3. Complete the following fields as needed:

Name

Specifies the name for the global collection.

Note: If your global collection name matches an existing global collection name, a warning appears. Click Yes to continue naming the global collection, or click No to enter another name for the global collection.

Owner

Specifies the user who is responsible for the global collection.

Description

(Optional) Specifies a description for the global collection.

Security String

(Optional) Specifies a security string expression to prevent certain users from viewing the contents of this global collection.

Note: For information about security string expressions, see the *Administrator Guide*.

Landscapes

(Optional) Changes the default landscape setting where this global collection is created.

Click OK.

The global collection of static members is created and appears in the Navigation panel under Global Collections.

More information:

[Provision Access to Modeled Elements](#) (see page 26)

Add Static Members to a Global Collection

You can add static members to an existing global collection.

Follow these steps:

1. In any topology, do *one* of the following steps to designate the modeled elements that you want to add to a global collection:
 - **Single modeled element selection:** In the Navigation panel, right-click a modeled element and select Add To, Global Collection(s).
The Select Global Collections dialog opens.
Note: Alternatively, you can right-click a single modeled element in a topology view and can select Add To, Global Collection(s).
 - **Multiple modeled element selection:** To multiselect modeled elements in a topology view, take the following steps:
 - a. Press and hold the SHIFT key and individually select the modeled elements.
 - b. While the SHIFT key is pressed, right-click the last selected modeled element and select Add To, Global Collection(s).
The Select Global Collections dialog opens.
Note: Or you can multiselect one or more modeled elements in the List tab and can add them to an existing global collection.
2. Select the name of the global collection where you want to add the modeled elements, and click OK.

The static members are added to the global collection.

Remove Static Members from a Global Collection

You can remove static members from an existing collection.

Follow these steps:

1. In the Global Collections navigation tree, right-click the static member that you want to remove from the collection and click Remove.

A dialog prompts you to confirm the deletion.

2. Click Yes.

The static member is removed from the collection.

The Remove operation removes the element from the collection, but it does not destroy the modeled element. If the modeled element exists in other topologies, it continues to exist in those topologies. If the modeled element does not exist in any other topology, it is placed in the Lost and Found and later destroyed.

Note: If you attempt to remove a *dynamic* member from a collection, an error message appears. The error message informs you that the selected member was added through a search criterion. In this case, redefine the search criteria to remove the dynamic member.

Edit Dynamic Members in Existing Global Collections

You can edit dynamic members in existing global collections.

Follow these steps:

1. In the Explorer tab of the OneClick Navigation panel, navigate to the Global Collections node.

2. Right-click the collection and select Edit Global Collection.

The Edit Global Collection dialog opens.

3. Edit the following fields as needed:

Name

Specifies the name for the global collection.

Note: If your global collection name matches an existing global collection name, a warning appears. Click Yes to continue naming the global collection, or click No to enter another name for the global collection.

Owner

Specifies the user who is responsible for the global collection.

Description

(Optional) Specifies a description for the global collection.

Security String

(Optional) Specifies a security string expression to prevent certain users from viewing the contents of this global collection.

Note: For information about security string expressions, see the *Administrator Guide*.

4. Click Search Options to modify the search settings for this global collection.

Note: If the search criteria of a global collection has been identified as having the capacity to result in degraded performance, make an audit.

Sometimes you cannot mitigate the potential performance impact by changing the search criteria. In this case, change the global collection Update Options to only update the collection members at a scheduled time.

To schedule the membership update:

- Select the Schedule button.

Notes:

- Avoid scheduling multiple updates at the same time, because it increases the potential to affect CA Spectrum performance.
- In a DSS environment that encompasses multiple time zones, the scheduled update times are local to each of the SpectroSERVERs. Consider this behavior when scheduling updates for any global collection that spans multiple landscapes.

Important! If a schedule is applied at the same time that the search criteria are changed, some landscapes could be updated when the changes are committed. This behavior is a known anomaly. To avoid this behavior, apply a schedule and commit the change before you change the search criteria.

5. Click Landscapes to modify the landscapes for this global collection.
6. Click OK

Your changes are saved and the dialog closes.

More information:

[Provision Access to Modeled Elements](#) (see page 26)

[Auditing Global Collections](#) (see page 103)

Auditing Global Collections

Several reasons could prompt you to audit a global collection, for example:

- Performing general housekeeping (identifying which collections are no longer needed).
- Determining who changed attributes of a collection, such as the update interval, search criteria, or name.

The following events can be used to gain information about changes that were made to global collections:

Event 0x1a100

Generated when the name of a global collection is changed.

Event 0x1a110

Generated when the owner of a global collection is modified and indicates the OneClick user that has modified the attribute.

Event 0x1a101

Generated when the search criteria is modified. This event indicates when the search criteria was changed and what it was changed to.

Event 0x1a111

Generated when the update method of a global collection is changed. You can use this event to determine when the method was changed and what it was changed to.

Note: If the changes are made to the update method for the global collection using CLI, the 0x1a111 and 0x1a110 events are not generated.

- Determining the mitigation of performance impacting collections.

When the search criteria of a global collection has the capacity to result in degraded CA Spectrum performance, make an audit. Using the audit findings, determine how to mitigate best the impact on the performance.

Some indications that search criteria for a dynamic global collection could be resulting in degraded CA Spectrum performance include:

- A SpectroSERVER performance event of type 0x10f20 or 0x10f21 has been generated on the global collection model.
- OneClick becomes unresponsive or disconnects from the Tomcat server during the dynamic update for the global collection.

For either symptom, we recommend that you examine the global collection to determine whether it is necessary. If the collection is still needed, the next step is to look at the search criteria to determine if it can be made more efficient.

More information:

[Global Collection Search Recommendations](#) (see page 94)

Copy Annotations from One Global Collection to Another

You can copy annotations (text) from one global collection and paste them into another global collection.

Note: You must have administrative privileges to copy annotations.

Follow these steps:

1. Expand Global Collections in the Explorer tab and select the global collection from which you want to copy annotations.
2. Click the Topology tab in the Contents panel.

Topology information for the global collection appears.

3. Click  (Edit) in the Topology tab toolbar.

Note: You must have administrative privileges to put a topology in Edit mode.

4. Select all of the annotations you want to copy and click Copy in the Topology tab toolbar. To select several annotations at once, press the CTRL key and click each annotation.

Note: If you select all of the annotations, the relative placement of the annotations is preserved.

5. Select the global collection that you want to copy the annotations to from Global Collections, in the Explorer tab.
6. Click the Topology tab in the Contents panel.

7. Click  (Edit) in the Topology tab toolbar and then click Paste.

The annotations are copied to the global collection you selected.

Copy Models from One Global Collection to Another

You can copy models from one global collection and paste them into another global collection.

Note: You must have administrative privileges to copy models.

Follow these steps:

1. Expand Global Collections in the Explorer tab and select the global collection from which you want to copy models.
2. Click the List tab in the Contents panel.
A list of all of the models in the global collection appears.
3. Select all of the models you want to copy and click Copy in the List tab toolbar. To select several models at once, press the CTRL key and select each model. To select a group of models at once, press the Shift key and select one model in the list. Then, select another model in the list. All of the models between the two selected models are also selected.
4. Select the global collection that you want to copy the models to from Global Collections in the Explorer tab.
5. Click the List tab in the Contents panel and then click Paste in the List tab toolbar.
The models are copied to the global collection you selected.

Find the Global Collections for a Model

You can determine whether a model belongs to a global collection of dynamic members, static members, or both.

Follow these steps:

1. Select the model for which you want to view global collection information.
2. Click the Information tab in the Component Details panel and scroll down to the Global Collections Memberships subview.
3. Expand the Global Collections Memberships subview.
The Static Global Collection Memberships subview and the Dynamic Global Collection Memberships subview appear.
4. Expand either subview to review the global collections that the model belongs to. If the model does not belong to any global collections, the tables within these subviews are empty.

Create a Global Collection Hierarchy

If you intend to organize your global collections using folders, set up the OneClick Navigation panel with a Global Collection Hierarchy. In this Global Collection Hierarchy, you can create multiple levels of folders to represent previously defined global collections.

Follow these steps:

1. In the Explorer tab of the Navigation panel, right-click the Global Collections Hierarchy node and select New Folder.
2. Type a descriptive folder name in the New Folder dialog and click OK.
The folder appears in the Global Collection Hierarchy tree.
3. In the Global Collection Hierarchy tree, perform any of the following tasks:
 - **Build more top-level folders:** Repeat Steps 1 and 2.
 - **Build one or more subfolders:** To create a subfolder, right-click a top-level folder and select New Folder.
 - **Populate a folder at any level:** To populate a folder with one or more collections, follow these steps:
 - a. Right-click the folder and select Add Global Collections.
 - b. In the Select Global Collections dialog, select the name of the collection you want to add and click OK.

Note: The Select Global Collections dialog represents the list of collections that are previously created in the Global Collection topology.

Modeling Manually in the World Topology View

You can represent your network geographically by creating a [World topology view](#) (see page 17). In a World topology view, you can model several layers of container views to depict your network locations. For example, you can create container views of network infrastructure from a national or regional level all the way down to an individual room that contains network equipment.

How to Model Locations

When modeling multiple containers representing locations within your network infrastructure, it is recommended that you use this process:

Step 1: Create top-level location views: From the World topology node, you can begin depicting the top-level view of any network location by modeling one of the following top-level containers:

- Building
- Site
- Region
- Country

Step 2: Create one or more sublevel location views: Depending on the top-level container that is modeled, you can then depict one or more sublevel containers.

For this top-level container:	You can model any of these sub-containers:
Building	Floor, Room, or Section
Site	Building
Region	Building or Site
Country	Region, Building, or Site

Step 3: Populate a room container view with modeled devices: After you have created a room container view, you can populate that view with modeled devices.

Note: The best practice is to model devices in the Universe topology view and then copy and paste these devices to the World topology view. You can use the Create Model by IP or the Create Model by Type option to manually model devices in a World topology view. However, this alternative approach is not recommended, because the Universe topology views represent the connectivity views of your network.

Define a Top-Level or Sub-Level Location View

You can create a top-level and sublevel location view, as well as populate a room container view with modeled devices.

Follow these steps:

1. In the Explorer tab of the OneClick Navigation panel, do one of the following steps:
 - **To define a top-level location view:** Click the World topology node to display the World topology view in the Topology tab of the Contents panel.
 - **To define a sublevel location view:** Click one of the top-level view folders appearing under the World topology node. The World topology view for that folder appears in the Topology tab of the Contents panel.



2. Click  (Creates a new model by type) in the Topology tab toolbar.

The Select Model Type dialog opens.

3. In the Container tab of the Select Model Type dialog, select a container type that best describes the network location you are depicting. Click OK.

The Create Model of Type dialog opens.

4. Specify a name that best describes the network location in the Name field.
5. Specify a security string in the Security String field if you want to secure this view from certain users.

Note: For more information about securing views, see [Provision Access to Modeled Elements](#) (see page 26).

6. Click OK.

The named icon container appears in the top-level (or sublevel) view of the World topology. The named folder representing the container appears in the Navigation panel under the World topology node.

7. (Optional) Click the Edit mode button in the World topology view to move the container icon or annotate this view further.

To populate a Room container view with modeled devices:

1. Go to the Universe topology view and copy the modeled devices that you want to display in a World topology view.
2. In the World topology view, navigate to a room container-type view.

3. In the room container-type view, paste the modeled devices.
4. Move the modeled devices to the desired location within this view.

Note: You cannot paste the same model into two different room containers. For example, you try to paste these same models into a different room container. You are asked whether you want to move them to the new room container or to keep them in the original room container.

If you determine you want to model new devices directly in the World topology, you can use one of these options:

- The Create Model by IP option
- The Create Model by Type option

More information:

[Add a Device Using Create Model by Model Type](#) (see page 74)

[Add a Device Using Create Model by IP Address or Create Model by Host Name](#) (see page 75)

[Editing and Enhancing Topology Views](#) (see page 135)

Modeling Manually in the TopOrg Topology View

You can manually model your network in the TopOrg topology when you want to group infrastructure models by organizational units or by services. For example, you can create a TopOrg topology view that depicts devices that are essential for supporting a network service, such as email. You can also depict services by department or by individual responsibility.

Note: When you populate a [TopOrg topology view](#) (see page 18), copy the modeled elements from a Universe topology view and paste them into a TopOrg topology view. Or you can use the Create Model by IP or the Create Model by Type option to manually model devices in a TopOrg topology view. However, we do not recommend this alternative because the Universe topology views represent the connectivity views of your network.

How to Model Services in the TopOrg Topology

Within the TopOrg topology, you can create multiple levels of containers. These multiple container levels represent organizations or individuals responsible for tracking the performance of mission critical services in your IT infrastructure.

When modeling multiple organizational containers in the TopOrg topology, consider following this process:

1. **Create layers of ownership or responsibility:** Using the Model by Type dialog, depict one or more containers that represent a department, individual, customer, or enterprise that is:
 - Supported by a network service, or
 - Responsible for tracking the performance of a network service.
2. **Populate Service_Owns container with supporting devices:** Populate the Service_Owns container with the modeled devices supporting the network service. You can populate these containers by copying and pasting modeled devices from the Universe topology to the TopOrg topology. You can also populate these containers by defining new devices using the Create Model by IP Address dialog.

More information:

[Define Service-Related Organizational Views](#) (see page 110)

[Populate Service_Owns or Org_Owns Containers](#) (see page 111)

Define Service-Related Organizational Views

You can create organizational and Service_Owns containers as well as populate a service_owns type container with modeled devices.

Follow these steps:

1. In the Explorer tab of the OneClick Navigation panel, click the TopOrg topology node.

The TopOrg topology view displays in the Topology tab of the Contents panel.
2. Click  (Creates a new model by type) in the Topology tab toolbar.

The Select Model Type dialog opens.
3. Click the Containers tab, select a container type that best describes the organization you are depicting, and click OK.

The Create Model of Type dialog opens.
4. Specify a name that best describes the organization that is responsible for tracking the performance of the network services or that these services supports.
5. (Optional) Specify a security string in the Security String field when you want to [secure this view from certain users](#) (see page 26).

6. Click OK.

The named container icon appears in the top-level view of the TopOrg topology. The named folder representing the container appears in the Navigation panel under the TopOrg topology node.

7. Repeat Steps 2 through 6 for each organizational or Service_Owns container you want to depict in the TopOrg topology.

Populate Service_Owns or Org_Owns Containers

Populate the Service_Owns container with the modeled devices supporting the network service. You can populate these containers by copying and pasting modeled devices from the Universe topology to the TopOrg topology. You can also define new devices using the Create Model by IP Address dialog.

Follow these steps:

1. Go to the Universe topology.
2. Copy the modeled devices that you want to display in a TopOrg topology view.
3. Go to the TopOrg topology view.
4. Navigate to a Service_Owns container view.
5. Paste the modeled devices.

Note: To model new devices in the TopOrg topology, you can use either the Create Model by IP option or Create Model by Type option.

Using Favorites

The Favorites folder contains modeled elements that the user has tagged for easy reference.

In the Navigation panel Explorer tab, you can add any OneClick element below the landscape level to your Favorites folder by right-clicking the element and choosing Add To, Favorites. You can also add Global Collections to your favorites by right-clicking your Favorites folder and choosing Add Collections.

To remove an element from the Favorites folder, right-click the element within the Favorites folder and choose Remove.

Important! If you right-click the element within the Favorites folder and you choose Delete, the element is removed from the Favorites folder. Plus, some models could also be deleted permanently from the system.

You can create subfolders by right-clicking Favorites (or a subfolder within Favorites) and choosing New Folder. Use the right-click menu to cut, copy, paste, rename, and delete subfolders.

More information:

[Deleting from Favorites](#) (see page 112)

Deleting from Favorites

Consider the following behaviors when deleting elements from the Favorites folder:

- When deleting containers from the Favorites folder, container models are permanently deleted from the system. Any models within the container are sent to the Lost and Found.
- When deleting a global collection from the Favorites folder, the global collection is permanently deleted from the system. Any models within the global collection are removed from favorites but are not deleted from the system.
- When deleting a model from the Favorites folder, the model is permanently deleted from the system. The model is *not* sent to the Lost and Found.

Lost and Found Model Information Subview

The Lost and Found Model Information subview lets you clear unattached models that are stored in the Lost and Found repository. The unattached models are models that have been cut but not pasted, models that Discovery could not resolve, and so on.

To access the Lost and Found Model Information view, select LostFound in the Navigation panel and select the Information tab in the Component Detail panel.

The Lost and Found model information view includes the following options:

Automatic Model Destruction

Specifies whether models in the Lost and Found are destroyed at specified times.

Next Model Destruction Date and Time

Specifies the next scheduled date and time at which models in Lost and Found are destroyed when Automatic Model Destruction has been enabled. The value of the Model Destruction Interval determines this value.

Model Destruction Interval

The interval (in seconds) at which models in Lost and Found are destroyed when Automatic Model Destruction has been enabled.

Default: 24 hours

Chapter 4: Configuring Modeled Devices and Interfaces

This section contains the following topics:

[Device and Interface Threshold Settings](#) (see page 115)

[Update Device Interface and Connection Information](#) (see page 119)

[Redundant Connections Between CA Spectrum and Modeled Devices](#) (see page 128)

[Interface Reconfigurations](#) (see page 131)

[Primary IP Address Modification](#) (see page 131)

[IPv6 Information](#) (see page 134)

Device and Interface Threshold Settings

CA Spectrum includes several device and interface alarms configured with three variables to define the alarm and reset conditions:

Threshold Setting

Specifies the threshold setting above which an alarm condition can exist.

Reset Level

Specifies the reset level below which an existing threshold alarm condition is automatically cleared. CA Spectrum does not generate subsequent alarms for the parameter until the value falls below the reset level.

Allowed Threshold Violation Duration

Specifies how long the parameter can be greater than the threshold setting, in seconds, before CA Spectrum generates an alarm.

You can configure device thresholds so that an alarm is generated if a given threshold is exceeded for a certain duration.

Note: To configure device and interface threshold alarm settings, the Device Threshold attribute on the SpectroSERVER must be enabled. You can turn off individual device or interface thresholds by setting the threshold and reset level variables equal to zero.

Device Threshold Settings


You can configure device thresholds so that an alarm is generated if a value is exceeded for a certain duration. No further alarms are generated for a threshold violation until its monitored value falls below the reset threshold. To disable a threshold from generating alarms, set the threshold value to zero.

The device threshold settings available are ‘% CPU Utilization’ and ‘% Memory Utilization’. You can access these OneClick settings in the following ways:

- The Thresholds and Watches subview in the Information tab for the selected device.
- The Attribute Editor, Thresholds, Device Thresholds grouping.

When using the Attribute Editor to set the Device Threshold attributes, the Allowed Threshold Violation Duration attribute is not available, and the default value of 300 seconds (five minutes) is used.

You can view the source that is used to calculate the CPU and memory utilization for a device in the Thresholds and Watches, Thresholds subview in the Information tab of the Component Detail panel.



More information:

[Threshold Attributes](#) (see page 190)

[Calculate Normalized CPU Utilization](#) (see page 197)

[Calculate Normalized Memory Utilization](#) (see page 198)

[How CA Spectrum Calculates CPU and Memory Utilization](#) (see page 191)

[Normalized CPU Utilization Attributes](#) (see page 193)

[Normalized CPU Utilization Calculation Requirements](#) (see page 192)

[Normalized Memory Utilization Attributes](#) (see page 195)

[Normalized Memory Utilization Calculation Requirements](#) (see page 193)

Example: % CPU Utilization Default Settings

This example illustrates how the default % CPU Utilization Threshold settings work together to trigger an alarm. The example is illustrated in the following graphic. The default settings for this alarm setting are as follows:

- % CPU Utilization Threshold = 85%
- % CPU Utilization Reset = 70%
- Allow Threshold Violation Duration = 300 seconds

Using the default settings, when the device's % CPU Utilization parameter exceeds the threshold setting of 85% at time Y, CA Spectrum begins the 300 second Allowed Threshold Duration timer.


The % CPU Utilization does not fall below the reset value of 70% for the duration of the timer.

At time = Y+5 minutes, CA Spectrum triggers a % CPU Utilization alarm for the device.

CA Spectrum will not generate another % CPU Utilization alarm until this alarm is cleared manually or automatically.

At time = Y+12 minutes, the device's % CPU Utilization falls below the reset value of 70%. CA Spectrum clears the % CPU Utilization alarm for the device.

Illustration: % CPU Utilization Default Settings



Interface Threshold Settings

The following interface thresholds parameters are available:

% Utilization Threshold

Defines the level of port capacity used that triggers an alarm condition for a port.

Packet Rate Threshold (packets/sec)

Defines the number of packets per second that triggers an alarm condition for a port.

% Error Rate Threshold

Defines the error rate on a port that triggers an alarm condition.

% Discarded Threshold

Defines the percentage of discarded packets on a port that triggers an alarm condition.

Each of these attributes has a reset value and Allowed Threshold Violation Duration timer attribute setting. You can access these OneClick settings in the following ways:

- The Thresholds and Watches subview in the Information tab for the selected device interface.
- The Attribute Editor, Thresholds, Interface Thresholds grouping.

Note: See [Threshold Attributes](#) (see page 190) for information about accessing Device Threshold settings using the Attribute Editor.

Update Device Interface and Connection Information

CA Spectrum can perform automatic discovery and mapping of a device's interfaces and connections based on the following events and conditions:

- A change in the number of configured interfaces on a device.
- When a device sends a LINK up trap.
- When CA Spectrum reconfigures a modeled device.

OneClick administrators can also manually update this information about a modeled device. See the *Operator Guide* for information about viewing a device's interface, sub-interface, and connection information.

Automatic Updates of Device Interface and Connection Information

You can use the following attributes to configure CA Spectrum to automatically update interface and connection information about a device.

- [Automatically Reconfigure Interfaces](#) (see page 120)
- [Discover Connections After Link-Up Events](#) (see page 120)
- [Create Subinterfaces](#) (see page 121)
- [Discovery After Reconfigure](#) (see page 122)
- [Topologically Locate Model](#) (see page 122)

Automatically Reconfigure Interfaces

When this attribute is set to Yes, CA Spectrum monitors the device for a change in the number of configured interfaces. If it detects a change, CA Spectrum automatically updates the device model to reflect the interface changes. The updated interface information appears in the Interfaces view for the device.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

Discover Connections After Link-Up Events

CA Spectrum automatically discovers and maps a model's connections one poll interval after it receives a Link Up trap from a device when the Discover Connections After Link-Up Events attribute is set to Yes. This delay lets the device fully reconfigure its related SNMP tables before CA Spectrum reads them. The Poll Interval setting for the device appears in the Information view, SPECTRUM Modeling Information subview.

Special Considerations for Flapping Interfaces

A "flapping" interface is one that is constantly coming up and going back down, usually because of a problem on the device. When Device Discovery After Link-Up Events is set to 'Yes,' CA Spectrum excludes LinkUp traps from flapping interfaces. As a result, the stream of LinkUp traps from a flapping interface does not interfere with a LinkUp trap for another interface on the same device. The connection discovery action can run as expected.

When CA Spectrum detects a flapping interface, a minor alarm is generated on the related device. After a default interval of 10 minutes without receiving a trap from that interface, the alarm clears. The default settings that are used to identify and track flapping interfaces are configured using Event Rules associated with Events 0x220002, 0x220006. The default settings are as follows:

- An Event Sequence Rule on Event 0x220002, which generates Event 0x220006 if a LinkUp trap is received and is followed by a LinkDown trap from the same interface within 60 seconds.
- An Event Rate Window Rule on Event 0x220006, which generates Event 0x220007 if 15 0x220006 events are generated within 5 minutes (300 seconds). Event 0x220007 generates a Minor alarm on the device.
- An Event Pair Rule on Event 0x220006 which generates Event 0x220008 if Event 0x220006 is not followed by Event 0x220007 within 10 minutes (600 seconds). Event 0x220008 clears the minor alarm generated by Event 0x220007.

The default values generate an alarm after 15 LinkUp/Down trap pairs are received. Also by default, the alarm is cleared 10 minutes after the last LinkUp/Down Trap pair is received. You can modify these settings by defining flapping interface event thresholds. The applicable rules are specified in the <\$SPECROOT>/SS/CsVendor/IETF/EventDisp file.

Note: For information about manually editing the files, for information about CA Spectrum events and event rules, and for information about changing the event rules associated with events 0x220002 and 0x220006, see the *Event Configuration User Guide*.

More information:

[SNMP Communication Attributes](#) (see page 189)

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

Create Sub-Interfaces Attribute

When this attribute is set to Yes, and the modeled device supports RFC 1573, CA Spectrum models the sub-interfaces for the device. CA Spectrum differentiates between physical and logical interfaces. It creates sub-interfaces using the logical interface information that it gathered from the device. A sub-interface appears in the Interfaces tab for a device, nested beneath the logical interface where it is configured.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

Discovery After Reconfigure

When this attribute is set to Yes, CA Spectrum rediscovers device connections each time that the device model is reconfigured. Both an interface reconfiguration and a manual device reconfiguration trigger a rediscovery when this attribute is enabled.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

Topologically Relocate Model

When this attribute is set to Yes, CA Spectrum determines whether a device model must be moved to a different topology during a Discovery. CA Spectrum moves the device if necessary, based on updated connection mapping.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

Manually Updating Interface and Connection Information

You can manually initiate an interface reconfiguration and connection discovery using the following options in the Tools, Reconfiguration menu and the Reconfiguration subview on the Information tab:

- [Reconfigure Model](#) (see page 122)
- [Discover Connections](#) (see page 124)
- [Rediscover SNMP MIBs](#) (see page 124)
- [Rename Interface Models](#) (see page 125)
- [Reevaluate Model Name](#) (see page 126)
- [Reevaluate NCM Device Family](#) (see page 126)

About Reconfiguring Models

When you activate a Reconfigure Model action, CA Spectrum finds the interfaces on the device and updates the device interface modeling.

Note: Reconfigure Model actions do *not* change the model type. To change the model type, either run the NewMM.pl script or delete and remodel the device by IP address. For more information about running NewMM.pl, see the *Installation Guide*.

The following parameters are reevaluated during a Reconfigure Model action:

Device Type

Verifies the current Device Type attribute value.

Model Name

Checks the Model Naming Order setting on the VNM, and determines whether the device model requires a change.

Application Discovery

Performs a Reconfigure SNMP MIBs action.

Interface Discovery

Determines which interfaces exist on a device, and updates the device modeling as needed.

Normalized Source

Verifies the attribute to use for gathering device CPU and memory usage information.

Serial Number

Verifies the device serial number, if available, and updates the device model if necessary.

802.3ad Trunk Memberships

Checks to see if the device interfaces are members of a 802.3ad trunk.

NCM Device Family

Checks the Device Family value that Network Configuration Manager uses to groups devices by vendor.

More information:

[Manually Updating Interface and Connection Information](#) (see page 122)

Reconfigure a Model

You can reconfigure a model to update device interface information. When you activate a Reconfigure Model action, CA Spectrum finds the interfaces on the device and updates the device interface modeling.

Follow these steps:

1. Locate the model.
2. Right-click the model and select Reconfiguration, Reconfigure Model.
The Reconfigure Model dialog opens, showing the progress of the requested action.
3. Click OK.
The Reconfigure Model dialog closes. The model is reconfigured.

Discover Connections

When you run the Discover Connections command, CA Spectrum performs a Discovery on the selected device. Discovery data lets CA Spectrum update and remap the device model connection information. You can also use this functionality to discover connections in a LAN container.

Follow these steps:

1. Locate the container for which you want to discover connections in the Explorer tab.
2. Right-click the container, and select Reconfiguration, Discover Connections.
The Discover Connections dialog shows the progress of the requested action. If connections are successfully discovered, this dialog indicates success.
3. Click OK.
The Discover Connections dialog closes. The connections among the devices in the selected LAN container now appear as pipes in the Topology view.

More information:

[Manually Updating Interface and Connection Information](#) (see page 122)

Rediscover SNMP MIBs

When a device model is created, CA Spectrum automatically creates models for each of the major and minor applications the device supports. Click Reconfigure SNMP MIBs to retrieve application support information from the device. The application models for the device are updated with any changes.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

[Manually Updating Interface and Connection Information](#) (see page 122)

Rename Interface Models

Use this function to update a device's interface model names after changing the device's Interface Name Primary Suffix attribute or the Interface Name Secondary Suffix attribute. Using this command forces CA Spectrum to rename the interface models using the current values of both the primary and secondary suffixes for the interface model. Some of the suffix options include ifName, ifAlias, ifDescr, and ifIndex.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

[Manually Updating Interface and Connection Information](#) (see page 122)

[Interface Configuration Attributes](#) (see page 185)

Entity Table Interface Stacking

When modeling interfaces, CA Spectrum uses the information contained within the MIB II ifStackTable to determine their logical stacking. For example, in the case of a frame relay interface with DLCI sub-interfaces, CA Spectrum attempts to stack the interfaces using information in the ifStackTable.

If you set the use_if_entity_stacking (0x12a83) attribute to TRUE on a device model in the Attributes tab, CA Spectrum attempts to use information from RFC2737 (Entity MIB) to determine interface stacking if the ifStackTable method fails. If an interface does not support ifStackTable, but the interface *does* support the Entity MIB, CA Spectrum will attempt to stack the interface model using information in the entPhysicalTable.

Note: This is done on a case-by-case basis as some vendors do not implement the RFC2737 indexing scheme correctly, which can cause interfaces to be incorrectly stacked.

More information:

[Edit Attributes in the Attributes Tab](#) (see page 171)

[Examine the Same Attribute on Multiple Models](#) (see page 173)

[Attribute Edit Panel](#) (see page 176)

[Create User-Defined Attributes](#) (see page 178)

Reevaluate Model Name

Determines whether to change the device's model name based on the VNM Model Naming Order setting for the VNM managing the device. See SpectroSERVER Control Subview for information about the VNM Model Naming Order setting.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

[Manually Updating Interface and Connection Information](#) (see page 122)

Reevaluate NCM Device Family

Automatically places a device in the proper device family after a firmware upgrade. For example, if you have a Cisco device that appears in the CatOS family and you then upgrade this device with new firmware and it changes to CiscoIOS, the device does not switch its family automatically. Instead, you can update it using the Reconfigure menu.

More information:

[Automatic Updates of Device Interface and Connection Information](#) (see page 120)

[Manually Updating Interface and Connection Information](#) (see page 122)

Access Interface and Connection Update Controls

You can access to the interface and connection update controls described in this section as shown in the following table:

Attribute	Tools, Reconfiguration Menu	Reconfiguration	Attribute Editor
Automatically Reconfigure Interfaces		X	X
Discovery Connections After Link Up Events		X	X
Discovery After Reconfigure		X	X
Create Sub-Interfaces		X	
Topologically Relocate Model		X	X
Reconfigure Model	X	X	
Discover Connections	X	X	

Attribute	Tools, Reconfiguration Menu	Reconfiguration	Attribute Editor
Rediscover SNMP MIBs	X	X	
Rename Interface Models	X	X	
Reevaluate Model Name	X	X	
Reevaluate NCM Device Family	X		

Tools, Reconfiguration Menu

The Tools, Reconfiguration menu provides quick access to reconfiguration actions you can perform on a selected device model. You can also access this menu by right-clicking the device you want to reconfigure.

Reconfiguration Subview and Advanced Subview

The Reconfiguration subview provides access to attributes that control when CA Spectrum updates a device's interface, connection, and topology information. You can also manually reconfigure a device and discover a device's connections from this subview.

The Advanced section of the Reconfiguration subview provides access to individual model reconfiguration functions that occur as part of the Reconfigure Model function. In some cases, you may want to perform these actions separately instead of performing an overall Reconfigure Model action.

Attribute Editor

Use the Attribute Editor to access some of the interface and connection update parameters for many models or modeled devices.

More information:

[Interface Configuration Attributes](#) (see page 185)

[Change Management Attributes](#) (see page 184)

Redundant Connections Between CA Spectrum and Modeled Devices

If a modeled device is configured with a pool of IP addresses available to it for use in communicating on the network, CA Spectrum can use these IP addresses to create redundant connectivity with that device. If the redundancy feature is enabled on a device and CA Spectrum cannot reach the device using the designated primary address, CA Spectrum attempts to re-establish contact using the list of available IP addresses.

More information:

[Configuring Allowed/Non-Alarming Shared IP Addresses](#) (see page 169)
[Change Management Attributes](#) (see page 184)

Redundancy Preferred Addresses List

A device that supports the CA Spectrum redundancy feature has a Redundancy Preferred Addresses list containing the device's interface IP addresses that is created when the router is originally modeled. CA Spectrum uses this list in determining redundant connectivity to the device. Devices have a primary address that is determined when the device is modeled. The CA Spectrum modeling process includes loopback functionality. If the VNM is configured to use the loopback feature, the first valid loopback address detected for the device is used as device model's primary address.

Removal of Shared IP Addresses from Preferred List

CA Spectrum automatically removes the IP addresses that it detects as "shared" from a Redundancy Preferred Addresses list for a device model. CA Spectrum places these shared addresses in the Redundancy Excluded Addresses list. Such shared IP addresses are not used when CA Spectrum attempts to restore communication with a lost device that was using redundant IP addresses.

If you manually add a shared IP address to the Redundancy Preferred Addresses list, CA Spectrum automatically moves it back to the Redundancy Excluded Addresses list. You cannot see this list until you close and reopen the current view.

More information:

[Shared IP Detection and Alarming](#) (see page 168)

Device Primary Address

By default, the primary address for a device is the IP address assigned to the device for network communications. You can change the primary address in OneClick if the network address for the device changes. If CA Spectrum cannot contact the device using its primary address, it attempts to contact the device using the first IP address in the Redundancy Preferred Addresses List.

You can change a device's primary address and the IP addresses listed in the Redundancy Preferred Addresses list using the IP Redundancy subview.

More information:

[Loopback Interfaces and Discovery](#) (see page 66)

IP Redundancy Subview

The IP Redundancy subview displays the attributes and settings CA Spectrum uses to create and monitor redundant communication paths to the device. Access the IP Redundancy subview by selecting a device in either the Navigation panel, the List tab, or the Topology tab, and then selecting the Information tab in the Component Detail panel. The IP Redundancy subview appears in the Information tab.

Enable Redundancy

When this attribute is set to Yes, CA Spectrum uses the addresses in a modeled device's Redundancy Preferred Addresses list, if it exists, to contact a device when the primary address is not available.

Generate Redundancy Alarms

When this attribute is set to Yes, CA Spectrum generates an alarm when a device cannot be contacted using its primary address.

Select Preferred Redundant Addresses

The Redundancy Preferred Addresses list displays IP addresses that a device can use for communicating on the network. You can manually add or remove IP addresses to or from the Redundancy Preferred Addresses list for a device.

Follow these steps:

1. Select the device with a redundant IP address.
2. Click the Information tab and expand the IP Redundancy subview.

3. Click Configure next to Primary Address.
The Preferred Addresses dialog opens.
4. Click Add below the Redundancy Preferred Addresses list.
5. Enter the IP address that you want to add to the list in the Add IP Address dialog, and click OK.
The IP address now appears in the Redundancy Preferred Addresses list.
6. Click OK.
The changes to the device's Redundancy Preferred Addresses list are applied.

Exclude Redundant Addresses

The Redundancy Excluded Addresses list displays IP addresses that a device cannot use for communicating on the network. You can manually add or remove IP addresses to or from a device's Redundancy Excluded Addresses list.

Follow these steps:

1. Select the device for which you want to add an excluded IP address.
2. Click the Information tab and expand the IP Redundancy subview.
3. Click Configure next to Primary Address.
The Preferred Addresses dialog opens.
4. Click Add below the Redundancy Excluded Addresses list.
5. Enter the IP address that you want to add to the excluded list Add dialog and click OK.
The IP address now appears in the Redundancy Excluded Addresses list.

Note: You can also select an address in the Preferred list and move it to the Excluded list by clicking the single right arrow located between the two lists. Similarly, you can move an address from the Excluded list to the Preferred list by selecting the address in the Excluded list and clicking the single left arrow located between the two lists.

About Shared IPs in Device Communication

Since any of your shared IP addresses will already have been added to the Redundancy Excluded Addresses list, the CA Spectrum redundancy intelligence will never attempt to communicate with a device using a shared IP. CA Spectrum will also never assign a shared IP address to the PrimaryAddress attribute.

Also, if you try to write a new value to a device model's NetworkAddress or PrimaryAddress attribute, if it is a shared IP address, the new value is not written and a warning dialog opens.

More information:

[Shared IP Detection and Alarming](#) (see page 168)

Interface Reconfigurations

Device models can potentially reconfigure interfaces at every poll cycle. This reconfiguration increases CPU usage and generates SNMP traffic due to interface table changes or interface stack table changes.

To disable interface reconfiguring, do the following:

- Set the Use_If_Table_Last_Change (0x11f7f) attribute to FALSE to disable interface reconfiguring for an interface table.
- Set the Use_If_Stack_Last_Change (0x130bc) attribute to FALSE to disable interface reconfiguring for an interface stack table.

You can configure a device model to trigger an alarm if it is continually reconfiguring interfaces. By default, a device model triggers a minor alarm if any one of the following sequences occurs in a 31-minute timeframe:

- Six interface reconfigurations for an interface table
- Six interface reconfigurations for an interface stack table
- Both sequences.

Primary IP Address Modification

The primary IP address is the address CA Spectrum uses to communicate with a modeled device. You can change a modeled device's primary IP address. There are three ways to change a device's primary IP address:

- [Change a device's primary IP address in the device's preferred address list](#) (see page 132).
- [Change a device's primary IP address to an interface's primary IP address](#) (see page 132).
- [Change a device's primary IP address to an interface's secondary IP address](#) (see page 133).

Change the Primary IP Address for a Device in the Preferred Address List

You can change the primary IP address for a modeled device. If you know the IP address that you want to use to contact the device, you can change the address in the preferred address list for that device.

Follow these steps:

1. In the Explorer tab, select the device whose primary IP address you want to change.
2. Click the Information tab and expand the IP Redundancy subview.
3. Click Configure next to Primary Address.

The Preferred Addresses dialog opens. The current primary address for the device appears in the Primary Address field.

4. Select an IP address from the Redundancy Preferred Addresses list on the left side of the dialog, and click Primary.

The IP address you selected appears in the Primary Address field. The IP address that had originally been the primary address now appears in the Redundancy Preferred Addresses list.

5. Click OK.

The change to the device's primary address is applied and the Preferred Addresses dialog is closed.

More information:

[Primary IP Address Modification](#) (see page 131)

Change the Device IP Address to an Interface Address

You can change the primary IP address of a modeled device to an interface address. In some situations, you do not know the IP address to contact a device, but you do know the device interface that you want to use. You can change the primary IP address for that device by selecting an interface primary IP address.

Note: For more information about interfaces, see the *Operator Guide*.

Follow these steps:

1. In the Explorer tab, select the device whose primary IP address you want to change.
2. Click the Interfaces tab in the Component Detail panel.

Interface information for the selected device opens in table format.

3. Right-click the interface you want to use to contact the device, click **Configure Primary Address for Device**, and then click **Use Interface IP As Primary Address** to set the interface's primary IP address as the device's primary IP address.

A confirmation dialog opens.

4. Click **Yes**.

The change to the device's primary address is applied.

Note: You cannot select an interface that does not have a primary IP address configured for it (for example, if the IP Address column in the interfaces table is blank). If you try to select the interface, an error message appears.

More information:

[Primary IP Address Modification](#) (see page 131)

Change the Primary IP Address for a Device to use an Interface Secondary IP Address

You can change the primary IP address of a modeled device. In some cases, you do not know the IP address that you want to use to contact a device, but you do know the particular interface. You can change the primary IP address for a device by selecting an interface and then selecting one of the interface secondary IP addresses.

Note: For more information about interfaces, see the *Operator Guide*.

Follow these steps:

1. In the Explorer tab, select the device whose primary IP address you want to change.
2. Click the Interfaces tab in the Component Detail panel.

Interface information for the selected device opens in table format.

3. Right-click the interface that you want to use to contact the device.
4. Click **Configure Primary Address for Device**.
5. Click **Use Secondary IP** to set one of the interface secondary IP addresses as the primary IP address for the device.

The Interface IP Mask Table opens.

6. Select an IP address to use from the list of secondary IP addresses and click **Use as Primary Address for Device**.

A confirmation dialog opens.

7. Click Yes.

The change to the device's primary address is applied.

Note: You cannot select an interface that does not have any IP addresses configured for it (for example, if the Secondary IPs and IP Address columns in the interfaces table are blank). If you try to select the interface, an error message appears.

More information:

[Primary IP Address Modification](#) (see page 131)

IPv6 Information

You can view IPv6 information for devices that support RFC2465 and RFC2452 MIBs in the Information tab of the Component Detail panel. The following tables provide specific information:

- IPv6 Interface Configuration Table
- IPv6 Routing Table
- IPv6 Address Table

Chapter 5: Editing and Enhancing Topology Views

This section contains the following topics:

[Topology Edit Mode](#) (see page 135)


[Set Topology View Edit Mode Preferences](#) (see page 137)

[Modifying Topology Views](#) (see page 138)

Topology Edit Mode

Topology Edit mode refers to the condition you place a topology view in when you want to edit its appearance. When you place a view into Edit mode, you automatically prevent other users from editing that view. Once in Edit mode you can use its drawing tools to draw rectangles, ellipses, lines, or text boxes. After you create these items in a view, you can later apply styles or colors to them.

The following image shows an example of the Edit mode toolbar that appears when you enter Edit mode:



Access Edit Mode

If your user account has the required privileges, you can access Edit mode to modify the current topology view.

Follow these steps:

1. Click the Topology tab in the Contents panel.
The Topology view and the Topology toolbar open.

2. Click  (Edit) in the Topology tab toolbar.

The Edit mode toolbar opens. The topology view is locked to prevent other users from editing this view.

3. Modify the topology view.
4. Click Save.

Your changes are saved, and you exit Edit mode.

More information:

[Resize Model Icons](#) (see page 139)

[Add Shapes, Lines, or Text to a View](#) (see page 140)

[Change Shapes, Lines, and Text Characteristics](#) (see page 141)

[Modify the Topology Background](#) (see page 143)

[Group Items in a View](#) (see page 144)

[Ungroup Items in a View](#) (see page 145)

[Send Items to the Back](#) (see page 145)

[Bring Items to the Front](#) (see page 145)

Edit Mode Toolbar

The following table describes the editing tools that you can access and use from the Edit mode toolbar.

Tool	Description
	Move Tool: Moves modeled elements in a view.
	Rectangle Tool: Draws rectangles. Click and hold the rectangle button to access additional tools.
	Ellipse Tool: Draws ellipses. Click and hold the Ellipse Tool button to access additional tools.
	Line Tool: Draws lines.
	Text Box: Creates a text box used to enter text.
	Font Properties: Opens the Select Font dialog for the selected text annotation. Choose a font family, style, and size from the respective columns in the Select Font dialog. You can also choose the text foreground and background color and whether to show the text background.

Tool	Description
	Shape Color: Opens the Select Shape Color dialog for the selected annotation. Select a shape color in the Select Color dialog.
	Line Weight: Sets the line weight for lines, ellipses, and rectangles.
	Line Pattern: Sets the style for lines, ellipses, and rectangles.
	Group: Group selected modeled elements in a view.
	Ungroup: Ungroup selected modeled elements in a view.
	Background Editor: Changes topology background characteristics (edit mode grid, grid spacing and color, background color, image, and size).
Arrangement	Arrangement drop-down list: Contains the following options for arranging the elements in the topology: Radial, Tree, or Manual.

Set Topology View Edit Mode Preferences

You can set preferences to specify how you want Edit mode to behave.

To set topology view Edit mode preferences

1. Click View, Preferences.
The Set Preferences dialog opens.
2. Expand the Topology Tab folder in the Name column.
3. Click any of the following options to make changes:

Annotation Font

Specifies the default font settings for topology annotation text. You can modify font, style, size, and background and foreground colors.

Grid Properties

Specifies the following settings for the grid that can appear in the Topology tab in Edit mode:

- **Show grid:** Set the size of the grid squares using the value displayed using the Show grid option. Decreasing the value decreases the size of the grid squares, while increasing the value increases the size of the grid squares.
- **Snap to grid:** Enables snap-to-grid while the topology view is in Edit mode, making it easier to align modeled device icons in the topology view.

Initial Zoom

Specifies the zoom behavior for topology views when they are first shown.

Show Pipe Label

Specifies whether you want to show pipe labels in the topology view.

4. Click OK.

Your changes are saved and the Set Preferences dialog closes.


Note: You can also set these preferences in the Background Editor dialog.

Modifying Topology Views

You can modify the appearance of any topology view by using the Edit feature in the view. Some of the enhancements you can make include the following:

- Change the background characteristics of a view.
- Add lines, rectangles, or ellipses to a view.
- Change the placement of modeled elements in a view.
- Change font characteristics in a view.

The following image shows an example of an enhanced topology view:



Multi-User Considerations

Be aware that OneClick topology view enhancements are shared across all users. Also, when you edit a view using the Edit mode button in a Topology tab toolbar, CA Spectrum automatically prevents all other users from editing that view until you have finished.

Resize Model Icons

You can resize model icons displayed in a topology view.

Follow these steps:

1. Switch to Edit mode, as described in [Access Edit Mode](#) (see page 135).
2. Select the model icon that you want to resize in a topology view.
A green box appears around the icon.
3. Click one of the corners of the green box, and drag the icon to resize it.
The icon size changes proportionally.
4. Deselect the icon when it reaches the desired size.
5. Click Save.

The model icon is resized, and you exit Edit mode.

Add Shapes, Lines, or Text to a View

You can add rectangles, ellipses, lines, or text in a topology view.

Add a rectangle or ellipse to a topology view in Edit mode.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).
2. Click  (Rectangle Tool) or  (Ellipse Tool) in the Edit mode toolbar and select the desired style for the shape from the menu.

The pointer changes from an arrow to a crosshair when hovering over the background area of the topology view.
3. In the desired location, click and drag the pointer starting at the upper left corner of the shape and ending at the lower right-hand corner of the shape.
4. Release the mouse button.

The shape you created now appears in the view, behind any existing models or pipes.

Add a line to a view in Edit mode.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).
2. Click  (Line Tool) in the Edit mode toolbar.

The pointer changes from an arrow to a crosshair.
3. In the topology view, click and drag the pointer to draw the line.
4. Release the mouse button.

The line appears in the view, behind any existing models or pipes.

Add text to a view.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).
2. Click  (Text Box) in the Edit mode toolbar.

The pointer changes from an arrow to a crosshair pointer.
3. Click in the location where you want the text box to begin.

A text box appears.

4. Type the text in the text box.
5. Click outside the text box to exit.

The text boundaries of the text box disappear. The text is placed in the background of the topology view.

Change Shapes, Lines, and Text Characteristics

You can apply different font properties to text, colors to shapes, or line weights to lines.

Note: We recommend setting properties for shapes, lines, or text before adding (or drawing) these elements in a view as a best practice.

You can change text font properties in a topology view.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).

2. Click the text you want to change and click  (Font Properties) in the Edit mode toolbar.

The Select Font dialog opens.

3. Select the desired font family, style, size, foreground color, and background color.

The preview pane shows the font properties that you selected.

4. Click OK.

The font properties are applied to the selected text.

You can apply color to shapes or lines in a topology view.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).

2. Click the shape or line that you want to edit and click  (Shape Color) in the Edit mode toolbar.

The Select Color dialog opens.

3. Modify the color settings by clicking each of the tabs in the Select Color dialog:

Swatches

Specifies the color of the shape or line. Select a color from the palette. A preview of the selected color appears at the bottom of the dialog. If you selected and previewed multiple colors, the colors you have chosen appear in the Recent color grid for re-selection.

HSB

Specifies the Hue, Saturation, and Brightness settings associated with standard color selected from the Swatches tab and shown in the preview at the bottom of the dialog. Use the slider to increase or decrease the settings associated with Red, Green, and Blue colors. Or, you can individually change the color settings associated with the Hue (H), Saturation (S), and Brightness (B).

Note: When you change the color settings in the HSB tab, the color settings in the RGB tab are updated respectively.

RGB

Specifies customization settings for a standard color chosen on the Swatches tab. Use the sliders to customize the standard color by adding more or less red, green, or blue.

4. Click OK.

The color settings are applied to the shape or line you selected.

You can also apply line weight and patterns to components of a topology view.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).

2. Select the line that you want to edit and click  (Line Weight) or  (Line Pattern) in the Edit mode toolbar.

A menu appears.


3. Click the desired line weight or line pattern.

The selected line weight or pattern is applied to the line.

Background Editor

Use the Background Editor dialog to modify the appearance of a topology view's background. You can modify a topology view's background color, add a background image, or change the size of the background. For example, you might want to change the background size of the topology view to create additional room for modeling network entities.

The following image shows the Background Editor dialog:



Modify the Topology Background

You can use the Background Editor to modify a topology background. You can change the background color, add a background image, or change the size of the background. For example, change the background size of the topology view to create additional room for modeling network entities.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).



2. Click (Background) in the Edit mode toolbar.

The Background Editor dialog opens.

3. Specify whether to change the color or image in the Background Type section.
4. Click Browse to preview colors or images.

The 'Select Topology Background Image' or 'Select Topology Background Color' dialog opens.

5. Select the desired image or color and click OK.
A preview of the selected item appears in the Background Editor dialog.
6. Click OK to apply the changes to the background.
The background view refreshes to reflect the changes.

You can use a similar procedure to change the background size.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).



2. Click (Background) in the Edit mode toolbar.

The Background Editor dialog opens.

3. Specify pixel values in the Width and Height fields.
4. Click OK.

The background view refreshes to reflect the changes.

Group Items in a View

You can group items in any OneClick topology view. Grouping items within a view lets you edit, move, copy, paste, or delete items as one group. One of the most common group operations you may perform within a topology view is to group text (annotations) with modeled devices.

To group items in a view

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).
2. Press and hold down the Shift key, and select the items you want to represent as a group.



3. Click (Group) in the Edit mode toolbar.

The selected items in the view are represented as a single group.

Ungroup Items in a View

Ungroup a set of grouped items within a view when you want to edit the items, or when you want to move individual items.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).
2. Click one of the grouped items to select the entire group.

3. Click  (Ungroup).

The items are ungrouped. You can now select individual items.

Send Items to the Back

You can send items to the back of a view so that other items within that view appear to be in front.

Follow these steps:

1. Switch to Edit mode, as described in [Access Edit Mode](#) (see page 135).
2. Right-click an item, and select Send to Back.

The item moves to the back of the view relative to other items in the view.

Bring Items to the Front

You can bring items to the front of a view relative to other items within that view.

Follow these steps:

1. Switch to Edit mode as described in [Access Edit Mode](#) (see page 135).
2. Right-click the item you want to move and select Bring to Front.

The item moves to the front of the view relative to other items in the view.

Chapter 6: Modifying Model Attributes

This section contains the following topics:

[Model Attributes](#) (see page 147)

[Attributes in the Information Tab](#) (see page 148)

[VNM Attributes in the Information Tab](#) (see page 149)

[Attributes Tab](#) (see page 170)

[OneClick Attribute Editor](#) (see page 174)

[Change Management Attributes](#) (see page 184)

[Interface Configuration Attributes](#) (see page 185)

[Maintenance Mode Attributes](#) (see page 186)

[Rollup Alarm Attributes](#) (see page 186)

[SNMP Communication Attributes](#) (see page 189)

[Threshold Attributes](#) (see page 190)

[How CA Spectrum Calculates CPU and Memory Utilization](#) (see page 191)

Model Attributes

Model attributes can be used to set values on models, set values directly on devices, turn CA Spectrum features on or off, configure CA Spectrum features, set default values in the CA Spectrum modeling catalog, and so on. You can modify an attribute which is associated with a device's MIB object, thus changing the device's value for that object without having to use the device's local management. For example, you can modify a device's contact details. You can modify the Maintenance or Hibernation mode attribute to control those features.

Important! Use caution when changing default settings for models; this can affect the overall performance of CA Spectrum. Additionally, attribute value changes you make to the selected models will affect the same values for similar models created in the future, and for any existing model if that model's type is using the default value.

You can review and modify model attributes as follows in OneClick:

- **Information tab:** Use the Information tab in the Component Detail panel to view and modify certain common attributes for a single model. See [Attributes in the Information Tab](#) (see page 148) for more information about viewing and setting attribute values for a single model from the Information tab.
- **Attributes tab:** Use the Attributes tab in the Component Detail panel to access every possible attribute for a selected model. You can also create custom views of attributes and modify each one as needed, depending on your access rights. See [Attributes Tab](#) (see page 170) for more information about viewing and setting attribute values for a single model using the Attributes tab.
- **Attribute Editor:** Use the Attribute Editor to modify non-list attributes for a model or subset of models as well as to modify default attribute values in the CA Spectrum modeling catalog. If you change any attribute values and apply them to the CA Spectrum modeling catalog, each device model that is subsequently created based on that model type will use the new attribute value. See [OneClick Attribute Editor](#) (see page 174) for information about using the Attribute Editor. You must have administrative read/write privileges for those models you want to configure with the Attribute Editor.

Attributes in the Information Tab

You can view and set attribute values for individual models using the Information tab. Subviews in the Information tab display grouped categories of information available for the model. The subviews available in the Information tab depend on the selected model.

The attribute values that appear in the Information tab for the selected model are a result of a combination of the following processes:

- Automated discovery and modeling
- Manual modeling
- Using the Attribute Editor
- Direct entry using OneClick
- Default CA Spectrum values

You can set some attribute values that appear in the subviews of the Information tab. Specific attributes that you can set depend on the model selected, privileges applied to model or model types, and other factors. You can change the values of attributes for which 'set' appears next to the attribute value.

More information:

[Model Attributes](#) (see page 147)

VNM Attributes in the Information Tab

In OneClick, you can view and set a variety of attributes for each Virtual Network Machine (VNM), or SpectroSERVER, in your CA Spectrum installation. The attributes available in the VNM Information view depend on the add-on applications installed as part of your CA Spectrum environment. The VNM attributes are grouped into subviews for specific applications and functionality. Most of the attributes have tooltips associated with them which describe the attribute.

General Information Subview

The General Information subview provides information about the VNM such as its network or IP address, condition, contact status, and when it was last polled successfully. With administrator privileges, you can set the VNM rollup alarm attributes. See [Rollup Alarm Settings](#) (see page 186) for more information. It also contains the following option:

Percent Models Activated

The percentage of models in the SpectroSERVER database that have been activated. The VNM icon will not change from its initial (blue) state until this value reaches 100%. This is useful to determine how close the SpectroSERVER is to becoming fully active after a restart. This value is also displayed in the message area of the CA Spectrum Control Panel.

CA Spectrum Modeling Information Subview

The CA Spectrum Modeling Information subview provides information about attributes such as SNMP community string, landscape, device type, and model type name.

Online Database Backup Subview

Use the settings available in this subview to configure online backups of the CA Spectrum database.

Automatic Backups

Specifies whether the CA Spectrum database is automatically backed up.

Default: Disabled.

Backup Interval

Specifies how often, in hours and minutes, the CA Spectrum database is automatically backed up.

New Backup Date & Time

Specifies the date and time of the next database backup.

Backup Compression

Specifies whether to compress the backup file using the default compression mode.

Default: Enabled.

Prefix for Backup File Name

Specifies the prefix used in the database backup file name. File names are appended with the date the backup occurred.

Backup Directory

Specifies the directory on the server where the backup files are written to. You must know the full path to the directory, as this is not a browse function.

Minimum Required Disk Space (MB)

Specifies the amount of free disk space that must exist on the server for a backup to start.

Note: You can initiate an online backup immediately by clicking Begin Backup Now.

SpectroSERVER Control Subview

The SpectroSERVER Control subview lets you configure various aspects of each of your local landscapes through various attributes and settings. It also contains the following views:

- [Alarm Information Subview](#) (see page 154)
- [Event Log Information Subview](#) (see page 155)
- [Statistics Log Information Subview](#) (see page 155)
- [Thread Log Information Subview](#) (see page 156)

The attributes and settings available in the SpectroSERVER Control subview include the following:

Device Thresholds

Set the Device Thresholds attribute to Enabled to activate the threshold functionality on devices supporting threshold. Each threshold values setting must also be set to a non-zero value for the threshold to be active.

Default: Enabled

Unmanaged Trap Handling

Specifies whether CA Spectrum processes “unmanaged” traps. Unmanaged traps are traps that come from devices which were not modeled in CA Spectrum. By default, the SpectroSERVER creates event records for any “unmanaged” traps it receives. As long as this setting is enabled, SpectroSERVER processes these unmanaged traps just as it processes traps from modeled devices; that is, until a trap “storm” occurs (as defined by the Trap Storm Rate and Trap Storm Length attributes).

The processing of unmanaged traps not only lets the network administrator know about unmodeled devices that may need to be modeled, but also allows monitoring of overall trap traffic. And it provides troubleshooting capabilities when traps are not mapped correctly. However, unmanaged trap handling can place a significant performance burden on the event logging and the Archive Manager. Depending on your priorities, you can use this setting to disable unmanaged trap handling entirely, or you can leave it enabled but limit it through the trap storm rate and length settings. Remember though that these settings also govern trap processing for modeled devices as well.

Note: Currently, only VNM models and EventAdmin models (created by users of the Southbound Gateway Toolkit) offer views that let you adjust these settings. For most device models, however, you can use the Attributes tab to create a custom view where you can adjust the default trap storm rate and length settings for that model. For more information about trap storm detection, see [How Trap Storm Detection Works](#) (see page 156).

Default: Enabled

Enable Trap Director

Lets you enable Trap Director when you want a given SpectroSERVER to forward incoming traps to models on remote landscapes in a distributed SpectroSERVER environment.

Default: Disabled

Auto Connects

Specifies whether CA Spectrum attempts to resolve the port connections when a pipe is created between two device models. This functionality will use the options that are enabled in the AutoDiscovery Control subview to resolve the port connections. Disabling Auto Connects can improve CA Spectrum performance if your modeled network contains management modules that support non-standard MIBs.

Default: Enabled

Copy Users When Copying Group

If the Copy Users when Copying Group attribute is set to Yes, whenever you copy a group or a user in a group from one landscape to another, the group and all users in the group are copied as well.

Default: Enabled

Log When Device Cannot Be Contacted

Specifies whether to continue logging attribute values (such as contact status) for models that have lost primary management contact with the devices they represent. In most cases, this is undesirable since it results in extra traffic to a part of the network where there may already be a problem. Hence, this option is disabled by default and logging is automatically suspended for a device when contact is lost.

VLAN Configuration

Specifies whether Virtual Local Area Networks (VLANs) are modeled for networks on this VNM.

Default: Disabled

Server Polling

Stops SpectroSERVER from polling the devices it is managing on the network. When SpectroSERVER polling is stopped, the VNM icon displays a gray condition status but no alarm will be generated. To restart SpectroSERVER polling of models, click Start.

Minimum Disk Space (kBytes)

Specifies the minimum amount of free disk space in kilobytes that must exist on the partition that the SpectroSERVER starts from for the SpectroSERVER to start. When the available space is less than this amount, a shutdown message is generated and the SpectroSERVER shuts down.

Default: 2000

Use Fully Qualified Host Name

Specifies whether the domain name is included with the host name when the Name Service selection is placed first in the Model Naming Order list. For example, if you select Yes here, the model's icon would be created with a fully qualified name such as myhost.ca.com. If you select No here, the model's icon would be created without a fully qualified name such as myhost. This only applies when you use the device name returned from the operating system.

Default: Yes

Model Naming Order

Specifies the order of the list of sources used by CA Spectrum to create model names for new models. If the first source at the top of the list is not available for a device, CA Spectrum attempts to use the next source in the list. The default order is as follows, with the top source being the first in priority:

- SysName
- IP Address
- Name Service

After changing the model naming order, click 'Reevaluate All Model Names' to have CA Spectrum run through all the models in the database and rename each one using the new model naming order.

The following additional scenarios will trigger the device model name to be reevaluated using the current model name selection. It will not reevaluate based on a new model name selection:

- If the IP address of the device changes and the model naming is based on IP Address or Name Service
- If the Reconfiguration, Reevaluate Model Name(s) action is manually applied
- If the Reconfiguration, Reconfigure Model action is manually applied

Note: If you do not want a specific device model name to be changed, set the value of the model's LOCK_MODEL_NAME (0x12a52) attribute to TRUE. This attribute locks the model name value so that it will not be changed.

Use Loopback

If Use Loopback is set to Yes, the SpectroSERVER will use the loopback interface as a primary agent address.

Default: No

Loopback if Description

Enter a string in this field to identify a preferred loopback interface for CA Spectrum to use when modeling the device. CA Spectrum compares the string entered with the if_descr entries in the device IFTABLE for loopback interfaces only. If a match is found, CA Spectrum uses that loopback interface when it models the device. If there is no match, or no value is specified, CA Spectrum chooses the loopback interface on the device with the lowest if_index value.

Update Event Configuration

Updates the SpectroSERVER with current alert and event mappings.

More information:

[Attributes Tab](#) (see page 170)

[Loopback Interfaces and Discovery](#) (see page 66)

[Trap Based Continuous Discovery Subview](#) (see page 160)

Alarm Information Subview

The Alarm Information subview provides the number of each type of generated alarm.

Active Alarms

Displays currently outstanding alarms by severity.

Total Active Alarms

Displays the sum of the outstanding alarms.

Total Alarms

Displays the break-down of the different types of alarms generated since the last server restart.

Note: Blue alarms that are caused by the creation of location or organization models are never cleared.

Total Alarms Generated

Displays the total number of alarms generated since the last server restart.

Event Log Information Subview

The Event Log Information subview provides information related to the event logs. This subview contains the following settings:

Events Generated

Indicates the total number of events generated since the last server restart.

Locally Stored Events

Indicates the number of event records currently held in the database. This field will read "0" unless the Archive Manager is shut down. This will serve as backup storage area for database records until the Archive Manager is restarted.

Events Purged

Indicates the number of event records written to the archive since the last server restart.

Max Log Size

Indicates the maximum number of event records held in the database. When this number is reached, records will be deleted.

Statistics Log Information Subview

The Statistics Log Information subview provides information related to the statistics logs. This subview contains the following settings:

Records Generated

Indicates the total number of statistic records generated since the last server restart.

Locally Stored Records

Indicates the number of statistic records currently held in the database. This field will read "0" unless the Archive Manager is shut down. This will serve as backup storage area for database records until the Archive Manager is restarted.

Records Purged

Indicates the number of statistic records written to the archive since the last server restart.

Max Log Size

Indicates the maximum number of statistic records held in the database. When this number is reached records will be deleted.

Thread Information Subview

The Thread Information subview provides information about the configuration and usage of threads. Comparing the In Use and Available columns for polling, logging, notification, and timer threads can help in determining if SpectroSERVER is running out of thread resources.

How Trap Storm Detection Works

The SpectroSERVER can block the processing of traps that are coming from managed and unmanaged devices when a threshold is reached. Excessive traps that are coming at a high rate can take down your SpectroSERVER and Archive Manager. You can enable the trap storm detection at your SpectroSERVER or at the level of a modeled device. When devices that are modeled in CA Spectrum send more than 20 traps per second, you must adjust `traps_per_sec_storm_threshold` so that trap storm detection does not limit the ability to receive traps.

You can enable the trap storm detection at any level by configuring the following two attributes. These attributes are available under the Attributes in the Component detail pane for the selected VNM model or for a selected device model:

`traps_per_sec_storm_threshold`

Defines the rate at which traps are received per second from a managed or unmanaged device. When this rate is sustained for the amount of time that is specified by the `TrapStormLength`, the SpectroSERVER stops the processing of traps from that unmanaged or managed device.

Default: 20 traps per second

`TrapStormLength`

Defines the time in seconds for which the `traps_per_sec_storm_threshold` value is sustained. SpectroSERVER considers it a trap storm and disables the processing of traps from that unmanaged or managed device.

Default: 5 seconds

When traps received from any device reach the configured thresholds, the SpectroSERVER identifies this rate as a trap storm. The SpectroSERVER stops handling traps from that device and traps from other devices are not blocked. SpectroSERVER trap storm detection logic is based on each IP address of an unmanaged or a managed device (trap source) that sends traps to SpectroSERVER. As a result, you can configure each device to send traps to the SpectroSERVER at the appropriate rate.

SpectroSERVER does not stop the processing of unmanaged traps when the overall trap storm rate from all the unmanaged devices exceeds the single trap storm threshold rate of an unmanaged device. As a result, you can configure each unmanaged device to send traps to the SpectroSERVER at the appropriate rate.

AutoDiscovery Control Subview

The attributes available in the AutoDiscovery Control subview affect actions that occur during Discovery and Modeling sessions. If you have a DSS environment, you must make any changes in these settings to all your SpectroSERVERs.

These parameters are applied when you are using the Discover LANs functionality available in a device model's Redundancy and Model Reconfiguration Options view, the Discover Connections functionality available from the right-click menu for a container model (LAN, Network, and so on), or the Auto Connects functionality used to resolve port connections when you manually draw a connection between two models. These parameters are also applied when you use the Discover Connections functionality with the Model by IP or New Model commands.

Each of these parameters is also available when you are selecting modeling options for Discovery. Parameters set in the AutoDiscovery modeling options override the default values for that AutoDiscovery.

Modeling and Protocol Options Subview

The modeling and protocol attributes affect how CA Spectrum discovers and models elements on a network using the following functionality:

- Discovering and modeling LAN functionality available when reconfiguring a device model.
- Discovering connections functionality for a container model (LAN, Network, and others).
- Auto Connects functionality used to resolve port connections when you manually create a connection between two models.
- Discover Connections functionality when creating a new model.

Create WA_Link Models

Creates a WA_Link model between the interfaces of two routers linked by a wide area connection. This occurs during layer 3 mapping. If this option is not selected, the two linked interfaces are directly connected without the WA_Link model. See [Wide Area Link Monitoring](#) (see page 221) for information about Wide Area Link models and how they are used.

Default: Yes

Create LANs (IP subnets)

Specifies whether CA Spectrum uses a LAN container to represent an IP Subnet. Discovery creates the LAN container during the Layer 3 mapping process for any router interface that routes to a local LAN.

Create Physical Addresses

When this option is enabled, a physical address model is created for any MAC address that is not associated with any modeled device but was heard by a switch. The layer 2 mapper attempts to find a connection for each address found. If a connection is found, a Fanout is created and the physical address is associated to it through Connects_To. If no connection is found the model is placed in Lost and Found. This option is not recommended.

Create 802.3 Fanout

If this parameter is set to Yes and if CA Spectrum cannot make an accurate connection among three or more interfaces, a Fanout model named "802.3_Segment" will be created and these interfaces will be connected to the Fanout model. If this parameter is set to No, a Fanout model will not be created for the interfaces that have unclear connection information, and therefore these interfaces will not be mapped. However, if there is a data relay device's interface among these interfaces, and all other interfaces are for end node devices, a Fanout model with name "Rpt_Segment" will be created.

Note: If you have 50 or more connections to a single Fanout model, consider changing this model to a Shared Media Link. The Shared Media Links must be modeled manually. These models can provide more control over fault management behavior when multiple connections are monitored. Unlike a Fanout model, Shared Media Links provide configurable thresholds for handling downstream connections that report problems. For example, a Fanout model reports a problem only when *all* downstream connections are down. However, a Shared Media Link can report the problem sooner, as when 60 percent of the downstream connections are down.

IP Address Tables

Discovery disables Layer 3 mapping and maps only the Layer 2 connections, when this option is disabled. In addition, when this option is disabled, Discovery automatically disables the IP Route Tables option, the Create WA_Link Models option, and the Create LANs (IP subnets) option.

Default: Yes

IP Route Tables

Specifies whether CA Spectrum will use the IP Address Table to map routers. This option is set to No by default because these tables can be very large and very time-consuming for CA Spectrum to read. When this option is enabled, CA Spectrum will not be able to map unnumbered IP interfaces (0.0.0.0).

Source Address Tables

If this is set to Yes, CA Spectrum will use the device's Source Address table when discovering connectivity information about this device.

Spanning Tree Tables

If this is set to Yes, CA Spectrum will use the device's Spanning Tree table when discovering connectivity information about this device.

Discovery Protocol Tables

Set the Discovery Protocol Tables attribute to Yes to allow CA Spectrum to map device connectivity using discovery protocol MIB information. Currently, the following discovery protocols are supported:

- Nortel Discovery Protocol
- Cisco Discovery Protocol
- Extreme Discovery Protocol
- Cabletron Discovery Protocol
- Alcatel Discovery Protocol
- Foundry Discovery Protocol
- Link Layer Discovery Protocol

Traffic Resolution

If the Traffic Resolution parameter is set to Yes, CA Spectrum will use network traffic data (ifInOctet and ifOutOctet statistics) to determine connections between interfaces, and in many cases eliminate the need for a Fanout model.

ARP Tables

When enabled, CA Spectrum uses the ARP table to determine pingable MAC addresses for the connectivity mapping.

ATM Protocols

If the ATM Protocols parameter is set to Yes, the ATM Discovery runs against all ATM switches in the SpectroSERVER database.

Default: No

SNMP Community Strings

Create, order, and delete community strings and profiles for SNMP v1, v2c, and v3, which are used, in order, when CA Spectrum attempts to access and model devices that were discovered using SNMP and for which no device community string was provided.

SNMP Ports

The SNMP Ports section lets you create, order, and delete the list of ports to use when accessing and modeling devices. To add port numbers to this list, click Add under the SNMP Ports field, enter the port number and click OK.

IP Exclusion List

A list of IP addresses or IP address ranges that will be ignored and which will not be modeled when devices are discovered.

More information:

[Wide Area Link Monitoring](#) (see page 221)

Trap Based Continuous Discovery Subview

Use the Trap Based Continuous Discovery subview to configure CA Spectrum to automatically create a device model when it receives an SNMP or syslog trap from a device not already modeled. When the SpectroSERVER receives an unmanaged trap, it asserts an event on the VNM model indicating that an unmanaged trap was received.

All models created using Trap Based Continuous Discovery are placed in the New Device Container. CA Spectrum places new models created by a scheduled continuous discovery or by an unmanaged trap into this container.

Unmanaged Trap Discovery

Set the Unmanaged Trap Discovery attribute to Yes to discover and model the source of an unmanaged trap using the IP address sent with the trap. This includes both SNMP and syslog traps from devices as well as Agent Log file matching traps. See SpectroSERVER Control Subview for information about how to enable unmanaged trap handling on the VNM.

New Devices In Maintenance

Set the New Devices In Maintenance attribute to Yes to have new device models created based on an unmanaged trap put into maintenance mode when they are discovered.

Create Pingables

Set the Create Pingables attribute to Yes to have CA Spectrum model devices that cannot be modeled using SNMP as type 'Pingable,' if the devices respond to a ping (ICMP) echo request.

Discover Connections

If the Discover Connections attribute is set to Yes, CA Spectrum attempts to discover and model the connections for devices discovered by Trap-based Continuous Discovery.

Debug Options Subview

The Debug Options subview lets you turn on the AutoDiscovery debugging functionality using the following settings:

Debug AutoDiscovery

Set the Debug AutoDiscovery attribute to 'On' to have CA Spectrum create a debug output file containing data on the status of the device modeling and mapping process for each Discovery session. These files are available at `<$$SPECROOT>/SS/ADiscDebug_<timestamp>`. The Debug AutoDiscovery option is useful when the discovery modeling or mapping process is hanging or when there are connectivity mapping issues. In these cases, the output file indicates where and on which devices any difficulties were encountered.

Note: In addition, you can debug a particular device's connectivity mapping process. To do this, set the Debug AutoDiscovery option to On.

When Discovery's modeling process is running, CA Spectrum prints out all connection information. This information includes the data collected from bridge tables, Proprietary Discovery Protocols, Spanning Tree tables, potential connections, errors encountered, and any additional, pertinent information related to the mapped devices.

Abort Discovery

The Abort Discovery button lets you stop and cancel a currently running AutoDiscovery.

Fault Isolation Subview

The Fault Isolation subview lets you configure various aspects of the CA Spectrum fault isolation functionality. For more information about this view, see [Fault Isolation Settings](#) (see page 201).

More information:

[Port Fault Correlation Options](#) (see page 204)

Live Pipes Subview

Live Pipes functionality lets you turn on port status monitoring for individual links and display the status of a link by using status color indicators. A link is a connection between two devices that CA Spectrum has resolved to the port level. For more information about Live Pipes and network fault management, see Live Pipes and Fault Management.

The Live Pipes attribute must be set to 'Enabled' on the VNM to turn on Live Pipes functionality on the VNM.

With administrator privileges, other attributes you can set in this view are Alarm Linked Ports, Suppress Linked Port Alarms, and Port Always Down Alarm Suppression.

Live Pipes and Global Collections in DSS Environments

In a DSS environment, the Live Pipes attribute must be set to the same value on all VNMs so that the Live Pipes functionality provides accurate link connection information in Global Collections. If Live Pipes is set to different values on VNMs in a DSS setup, the Live Pipes information in Global Collections will be unpredictable.

Alarm Management Subview

The Alarm Management subview lets you control some aspects of alarm management.

The AlarmMgmt model, which governs the Alarm Management subview, is a SpectroSERVER application. The AlarmMgmt model inherits the security string of the VNM model only if you have not independently changed the security string of the AlarmMgmt model.

For example, the security strings for the VNM and AlarmMgmt models are initially empty. You change the AlarmMgmt model security string to "Jack" and later, you change the security string for the VNM model to "Jill." The AlarmMgmt mode security string is not changed to "Jill."

AlarmMgmt model attributes are not distributed. Bring up the Alarm Management subview for each SpectroSERVER whose alarm management attribute values you want to change. Changing an attribute on one SpectroSERVER does not apply to any other SpectroSERVER.

Important! Displaying Initial and Suppressed alarms is not recommended in OneClick. These alarms can generate a significant volume of network traffic. CA Spectrum generates initial and suppressed alarms if the Disable Initial Alarms and Disable Suppressed Alarms settings for the Virtual Network Machine (VNM) managing your network are set to Yes.

The Alarm Management subview contains the following attribute settings:

Generate Alarm Events

Enables the generation of alarm change events; CA Spectrum creates events (viewable in the Events tab) for alarm changes based on alarm creation, updating, and clearing events.

Note: If the Generate Alarm Events option is disabled, you do not see Alarm History in the Alarms view.

Default: Yes

Add Events to Alarms

Controls whether alarm change events are added to each alarm. If disabled, alarm change events are not displayed under the Events tab of the Alarm view.

Default: No

Age Out Residual Alarms Only

Specifies whether only residual old alarms are cleared. Residual alarms are alarms that existed prior to a SpectroSERVER restart and have not been reverified. If enabled, CA Spectrum clears only residual alarms based on the Alarm Age Out timer setting.

Default: Yes

Alarm Age-Out Time (hours)

Defines how long an alarm can exist in CA Spectrum. Once an alarm has existed for the number of hours specified by this attribute, it is a candidate for automatic removal. To disable this functionality, set this attribute to zero (0).

Every hour, CA Spectrum checks the status of all alarms in the landscape and uses this option to determine whether alarms are cleared. Therefore, an alarm is not removed at the precise moment when its existence time has exceeded the time-out. An alarm can be, at most, an hour “overdue.”

Disable Initial Alarms

Specifies whether to generate an alarm if the condition of a model changes to Initial. Because these changes in the condition of a model can cause a flood of alarms, disabling this option can improve system performance.

Default: Yes

Note: If Initial, Suppressed, or Maintenance alarms are disabled and later enabled, you will not see these alarms in the Alarm view for existing models. Only alarms generated after this option is enabled appear in the view.

Disable Suppressed Alarms

Specifies whether to generate an alarm if the condition of a model changes to Suppressed. Because these changes in the condition of a model can cause a flood of alarms, disabling this option can improve system performance.

Default: Yes

Disable Maintenance Alarms

Specifies whether to generate an alarm if the condition of a model changes to Maintenance. Because these changes in the condition of a model can cause a flood of alarms, disabling this option can improve system performance.

Note: For more information about putting devices in Maintenance mode, see the *Operator Guide*.

Default: No

Auto UnAcknowledge On New Occurrence

Specifies whether to unacknowledge the new occurrence of an alarm.

Default: No

BGP Manager Subview

The BGP Manager subview lets you globally configure BGP peer session monitoring.

The BGP Manager subview contains the following attribute settings:

BGP Peer Session Monitoring

Monitors the status of the peer session on the BGP port at the polling interval of the port model's Polling_Interval Attribute value if this setting is enabled and the live pipe on the BGP peer session port is turned on. If you disable this option, an event of type 0x220018 is generated on the BGP downed port models to clear the BGP alarm.

Default: Disabled

BGP Peer Session Discovery Interval (minutes)

Indicates the interval for BGP peer session Discovery. If BGP Peer Session Monitoring is enabled, BGP Peer Session Discovery initially runs on each BGP device at SpectroSERVER startup and when a new BGP device is modeled. After, BGP Peer Session Discovery runs according to the interval you set.

Default: 24 hours

More information:

[BGP Peer Session Monitoring](#) (see page 84)

Network Configuration Manager Subview

The Network Configuration Manager subview provides information about Network Configuration Manager.

This subview contains the following setting:

Export Directory

Specifies the local directory to which you want to export configuration text files. If you want to export configuration text files to a network share, you must specify the UNC path to the directory. For example, \\Shared_Server\\Export\\ExportFiles.

TFTP Configuration Subview

The TFTP Configuration subview provides information about the Trivial File Transfer Protocol (TFTP). TFTP transfers configuration files.

This subview contains the following settings:

Default TFTP Directory

Specifies the TFTP server path.

TFTP Transfer Timeout (sec)

Specifies the maximum time (in seconds) for a data transfer to complete.

Default: 50 seconds, which means the data must be completely transferred within 50 seconds.

Note: For more information about the TFTP server, see the *Network Configuration Manager User Guide*.

FTP Configuration Subview

The FTP Configuration subview provides information about the File Transfer Protocol (FTP).

This subview contains the following settings:

FTP Username

Specifies the FTP server username.

FTP Password

Specifies the FTP server password.

Default FTP Directory

Specifies the FTP server path.

Note: For more information about the FTP server, see the *Network Configuration Manager User Guide*.

Thresholds And Watches Subview

You can create, configure, and administer watches in OneClick. View and configure watches from a table in the Thresholds And Watches subview.

Note: You can access the Thresholds and Watches subview from the Information tab for a model.

The Watches table displays information for each watch defined on that model. The Watch Status column displays the watch condition with color codes as follows:

Gray

Indicates that the watch is inactive. The watch is not currently running because it has not been activated.

Blue

Indicates the initial state of the watch. The watch is activated but has yet to run for the first time.

Green

Indicates that the watch is active and running without any violation.

Yellow

Indicates that the watch threshold is violated.

Red

Indicates that the watch failed to evaluate. The text explains the reason.

The toolbar buttons let you do the following:

- Activate
- Deactivate
- Create
- Edit
- Copy
- Delete
- Display watch information
- Print watch information
- Export the Watches table

Host Security Information Subview

When a client application connects to a SpectroSERVER, CA Spectrum reads the .hostrc file to obtain a list of valid hosts. If a host name from the .hostrc file does not resolve to a network address, you will receive a “Permission Denied” error message. In addition, an event and an alarm (Event00010e01, Prob00010e01) will be generated on the VNM indicating that there are unresolved host names.

To help you find the cause of this problem, the Host Security Information subview displays a list of resolved and unresolved host names.

Modeling Gateway Subview

You can view information about recent imports in a table in the Modeling Gateway subview.

The Modeling Gateway table displays information about recent imports. The number of import files listed is controlled by the Max Records field. The default value for the Max Records field is 30.

Note: For more information about the Modeling Gateway table, see the *Modeling Gateway Toolkit Guide*.

IP Services Subview

The IP Service subview provides information about VPN Manager and VPLS Manager. Further options are available depending upon the products you have installed.

Logical Connection Import Subview

The Logical Connection Import subview lets you create logical connections between virtual link models by importing a comma-delimited, ASCII file (text file or XML file) that defines the connections. You can define connections that include two ATM models or an ATM model and a Frame Relay model. Click the Import button to import a file.

Note: For more information about logical connections between virtual link models, see the *ATM Circuit Manager User Guide*.

Shared IP Detection and Alarming

The following settings control when CA Spectrum generates alarms for shared IP addresses.

Shared IP Alarming

Specifies whether shared IP alarming is enabled.

Default: Disabled

Note: Shared IP alarms will be cleared when you set the Shared IP Alarming attribute to Disabled.

Currently Shared IP Addresses

Specifies which IP addresses are currently considered “shared” in CA Spectrum.

Note: The IP addresses in the loopback subnet are displayed as shared addresses in the ‘Currently Shared IP Addresses’ list, however, no alarms will be triggered based on these addresses to help prevent multiple unnecessary alarms for a known and desired configuration.

Allowed Shared IP Addresses

Specifies which IP addresses can be shared in CA Spectrum. Click Add or Remove to modify this list as needed.

Note: You can modify a device model's NETWORK_ADDRESS (0x12d7f) attribute and PrimaryAddress (0x12d80) attribute if the device model's IP address is included in the Allowed Shared IP Addresses list on the VNM model.

More information:

[Removal of Shared IP Addresses from Preferred List](#) (see page 128)

[About Shared IPs in Device Communication](#) (see page 130)

Shared IP Alarms and Events

When CA Spectrum detects that two or more devices share one or more IP addresses, and you have configured CA Spectrum to generate alarms in this case, you will see an orange alarm on all device models that share the IP address or addresses. The event generated on each device will contain a list of all device models involved as well as a list of all shared IPs. The event will look similar to the following:

Device {X} of type {Y} has the following shared IP addresses:
<list of shared IPs and devices>

Since the detection of shared IP addresses is dependent on CA Spectrum device models, each time a new device model is created or destroyed, a new event containing updated data may need to be generated on the devices that have shared IPs.

No Unique IP Alarms and Events

If a device is found to contain no unique IP addresses, then a red alarm is asserted on it to notify you of this condition because no reliable communication or management may be made with that device. The event will look similar to the following:

Device {X} of type {Y} has no globally unique IP addresses. Each of the following addresses is shared with another device:

<list of shared IPs and devices>

Network Address Is Shared

If you manually create a device using a shared IP address as the Network_Address you will receive an event such as the following:

Device {X} of type {Y} has its Network Address set to an IP that is currently shared by multiple devices. No reliable communication or management may be made with this device. The shared IP {shared IP} is shared by the following other devices:

<list of other devices>

Configuring Allowed/Non-Alarming Shared IP Addresses

You can configure CA Spectrum with a list of IP addresses, IP address ranges, or subnets for which sharing between multiple devices is allowed. Populate the Allowed Shared IP Addresses list with the addresses to share. The IP addresses on this list do not generate alarms.

You can use OneClick to add or remove IP addresses, IP address ranges, or subnets from this list. Shared alarms are cleared when you add an IP address to the list. Adding an IP address causes the associated device to have no more shared IP addresses that generate alarms.

More information:

[Redundant Connections Between CA Spectrum and Modeled Devices](#) (see page 128)

Attributes Tab

The Attributes tab in the Component Detail panel provides access to all of a selected model's attributes. From the Attributes tab you can select one or more attributes related to the specific model or model type and review details, poll values, export values, and edit each attribute as needed, depending on your access rights. You can also use the Attributes tab to cycle through models, quickly checking the same attribute for each one, to review attribute flags by scanning the Flags information located at the bottom of the Attributes view, or to review values for list attributes individually or all at once, depending on your preference.

The following graphic shows an example of the Attributes tab. The attributes being shown in the Attributes tab belong to the model selected in the List tab:

Contents: cisco4500 of type Cisco4500

Alarms | Topology | List | Events | Information

Filter: [] Displaying 89 of 89

Condition	Name	Network Address	Secure Domain	Manufacturer	Model Class	MAC Address	Type
Normal	172.19.10.0	172.19.10.0			Network		LAN
Normal	172.19.30.0	172.19.30.0			Network		LAN
Normal	172.19.4.0	172.19.4.0			Network		LAN
Normal	172.19.5.0	172.19.5.0			Network		LAN
Normal	172.19.55.0	172.19.55.0			Network		LAN
Normal	172.19.57.0	172.19.57.0			Network		LAN
Normal	172.19.58.0	172.19.58.0			Network		LAN
Normal	172.19.59.0	172.19.59.0			Network		LAN
Normal	172.19.6.0	172.19.6.0			Network		LAN
Normal	172.19.7.0	172.19.7.0			Network		LAN
Normal	172.19.71.0	172.19.71.0			Network		LAN

Component Detail: 172.19.58.0 of type LAN

Information | Post Configuration | Root Cause | Interfaces | Performance | Dashboard | Alarms | Events | Attributes

Filter: [] Displaying 193 of 193

Name	ID	Type
cisco4500	0x122e0	Integer
HardErrorRate	0x11559	Integer
Condition	0x1000a	Integer
Condition_Value	0x1000b	Integer
Value_When_Yellow	0x1000c	Integer
Value_When_Orange	0x1000d	Integer
Value_When_Red	0x1000e	Integer
Composite_Condition	0x1000f	Integer
Yellow_Threshold	0x10010	Integer
Orange_Threshold	0x10011	Integer
Red_Threshold	0x10012	Integer
Rollup_Condition	0x10013	Integer
GlobeAutoPlace	0x12a94	Integer
GlobeEditCount	0x12a9c	Integer

Name	Value
Yellow_Threshold	3
Orange_Threshold	6
Red_Threshold	10

Flags: Database Memory Readable Writable External Polled Logged Shared Global

Instance ID: []

More information:

[Model Attributes](#) (see page 147)

Access Attributes from the Attributes Tab

You can access attributes from the Attributes tab and personalize your view of them as required.

Follow these steps:

1. Select the model whose attributes you want to view or edit.
2. Click the Attributes tab in the Component Detail panel.
A list of attributes appears in the left half of a split panel.
3. (Optional) Enter text in the Filter field at the top of the list to filter it.
4. Double-click each attribute that you want to display in a view.
Each attribute that you double-click appears in the right-side of the panel along with its value.
5. (Optional) Click a column header to sort the attributes as needed.
6. (Optional) Select an attribute in the right-side of the panel and click the left arrow button at the top of the panel to move the selected attribute back to the left-side of the panel when you no longer want to review it.

Edit Attributes in the Attributes Tab

You can edit attributes for a single model from the Attributes tab.

Follow these steps:

1. Select the model whose attributes you want to modify.
 2. Click the Attributes tab in the Component Detail panel.
The model's available attributes appear in the left side of the panel.
 3. Double-click each attribute that you want to edit.
Each selected attribute and its value appear in the right side of the panel.
 4. Double-click an attribute in the right side of the panel.
If the selected attribute can be modified, the Edit dialog opens.
 5. Clear the 'No Change' check box in the Edit dialog to enable editing.
 6. Modify the attribute as needed, and click OK.
The Attribute Edit Results dialog opens, indicating whether the attribute edit was successful.
- Note:** Click Undo in the Attribute Edit Results dialog to revert to the original attribute settings.

7. Click Close.

The attributes have been edited and the Attribute Edit Results dialog closes.

8. (Optional) Click Export to send the selected attributes and their values to a CSV file, a text file, or web page.

More information:

[Entity Table Interface Stacking](#) (see page 125)

Edit Multiple Attributes at Once in the Attributes Tab

You can edit multiple attributes simultaneously in the Attributes tab.

Follow these steps:

1. Select the model whose attribute values you want to modify.
2. Click the Attributes tab in the Component Detail panel.

The available attributes for this model appear in the left side of the panel.

3. Double-click each attribute that you want to edit.

Each selected attribute and its value appear in the right side of the panel.

4. Click the Edit button in the toolbar of the right panel.

The Edit dialog lists the attribute values that are available for editing from your selected list.

5. Clear the 'No Change' check box in the Edit dialog for each attribute, modify each attribute as required, and click OK.

The Attribute Edit Results dialog opens, indicating whether each edit operation was successful.

Note: Click Undo in the Attribute Edit Results dialog to revert to the original attribute settings.

6. Click Close.

The attributes have been modified. The Attribute Edit Results dialog closes.

7. (Optional) Click Export to send the selected attributes and their values to a CSV file, a text file, or web page.

Examine the Same Attribute on Multiple Models

From the Attributes tab, while you are in either the List tab or the Topology tab, you can select an attribute or multiple attributes and quickly view the attribute value on a number of models.

Use the List tab to examine the values of the same attributes for multiple models.

Follow these steps:

1. Select the first model for which you want to modify attribute values from the List tab in the Contents panel.

2. Click the Attributes tab in the Component Detail panel.

The model's available attributes appear in the left side of the panel.

3. Double-click each attribute whose values you want to review.

Each selected attribute and its value appear in the right-side of the panel.

4. Press the down or up arrow on your keyboard to review the same attributes for a different model.

The Attributes view refreshes to display the values for the same attributes on the selected model.

Note: If you move to a different container, these attributes remain selected. The attributes stay in the Attributes tab until you exit OneClick.

More information:

[Entity Table Interface Stacking](#) (see page 125)

View List Attribute Values

You can use the Attributes tab to review the values of a model's list attributes.

To view the value for a particular instance of a list attribute

1. Select the model for which you want to view attribute values.

2. Click the Attributes tab in the Component Detail panel.

The model's available attributes appear in the left side of the panel.

3. Double-click the list attribute for which you want to review values.

The list attribute and the value of the first instance in the list appear in the right-side of the panel.

4. Do *one* of the following to review the list attribute's values:

- In the Instance ID field at the bottom of the Attributes tab, type the OID of the particular value you want to view, and press Enter.

Note: The Instance ID applies to the list attributes in the right-side panel. If you place a new list attribute in the right-side panel, the value displayed in the Value column corresponds to the OID specified in the Instance ID field.

- Click the table link in the Value column to open a dialog which displays the instances and values for the list attribute.

Note: You can perform a number of actions from this table including the following: refresh values, print values, and export values.

Update Attribute Values

The values of selected attributes are not dynamically updated. They reflect the value returned as of the SpectroSERVER poll prior to their selection.

To update attribute values, click Refresh in the toolbar in the right panel of the Attributes tab.

The values of all the attributes you originally selected from the left panel refresh to display any new values.

OneClick Attribute Editor

The Attribute Editor is an advanced CA Spectrum utility used to configure management 'policies' that govern how CA Spectrum manages network devices and their components. It is best suited for performing bulk attribute changes on multiple devices models.

You can change attribute values for one or more selected models in a view. The Attribute Editor dialog groups attributes into categories. You can edit the default settings provided within these categories, or you can define additional attributes to edit within the User Defined category.

More information:

[Model Attributes](#) (see page 147)

Open Attribute Editor

You can open the Attribute Editor in OneClick by right-clicking any model and selecting Tools, Utilities, Attribute Editor.

You can also launch the Attribute Editor from anywhere in OneClick where you can select a model, including the List tab, the Explorer tab, Interfaces tab in Component Details panel, the Locator Results tab, or from the Tools menu.


Open Attribute Editor with Device Context

You can open the Attribute Editor within the context of selected models. Select a model or multiple models in the List tab, the Explorer tab, or the Interfaces tab. Then right-click and select Utilities, Attribute Editor. The Attribute Editor opens with the context of the selected models.

Any changes that you make apply to the selected models. If you select Set as Default, the changes are applied to the CA Spectrum modeling catalog for the selected models.

Attribute Editor Dialog

The Attribute Editor includes a right and left panel. The left panel groups attributes in a tree display. The right panel provides an editing area to view current attribute values and make changes.



Task-Oriented Attribute Groupings

The left panel of the Attribute Editor provides attributes grouped by tasks you perform on devices and by categories of attribute types. The SNMP Communications folder groups attributes related to tuning SNMP communications between the SpectroSERVER and a device.

Filter Attribute Categories

You can type text in the Filter text box to locate attributes in the attribute categories in the Attribute Editor left panel. When you type text in the Filter text box, attribute categories that do not contain an attribute that matches the filter appear as bullets, becoming inaccessible. All attribute categories that contain an attribute that matches the filter appear as folders that you can expand and you can select the subcategories within to move into the right panel for editing.

For example, if you want to find attributes related to alarms, type **alarm** in the Filter text box.

Attribute Edit Panel

To edit attribute values, you must select the attribute category from the left panel, then click the right facing arrow to place the attributes in the right side editing panel. When an attribute is placed in the edit panel, it appears in *italics* in the attribute list in the left panel.

Tool tips describing an attribute are available for some attributes when they appear in the right panel for editing. Place the cursor over the attribute name and if a tool tip is available, it appears.

- **No Change:** The No Change setting appears to the left of most attribute input fields. When No Change is selected, the input value, if any, is not written when you click Apply or OK. When you make changes to the attribute value by selecting a value or clicking in the input field, No Change is automatically deselected. Either deselecting No Change or clicking in the input field makes the attribute value editable.
- **Set as Default:** This option appears to the right of most attribute input fields in the Editor panel. If you select Set as Default, the value is written to the model types in the CA Spectrum modeling catalog when you click Apply or OK. All future models that use these model types inherit the new value.

Important! Altering the default model-type value of an attribute changes the CA Spectrum modeling catalog. Any existing models that use that current default value and any models created in the future will use the new default value. This means changing the default value could affect models that are not explicitly selected. Changes to models not explicitly selected may not take effect until after the SpectroSERVER restarts.

- **Load Attribute Values:** You view the current value for a set of selected attributes when you have launched the Attribute Editor in the context of a specific model. After populating the edit panel, click Load Attribute Values to view the current values for the attributes. If the selected models do not use or have a value for an attribute, no value displays when you click Load Attribute Values.

If you have launched Attribute Editor in the context of multiple model types, they may have different values for the same attribute. If this occurs, the Select Model dialog opens when you click Load Attribute Values. Select the model you want to load attribute values for, and click OK.

When you click Apply or OK in the Attribute Editor dialog, CA Spectrum attempts to write the new attribute values and displays the Attribute Edit Results dialog.

More information:

[Entity Table Interface Stacking](#) (see page 125)

Attribute Edit Results Dialog

The results of attribute value changes appear in the Attribute Edit Results dialog. Each item in the table represents the result of a single attribute written to a model. The Result column indicates whether the write operation succeeded or failed. The Old Value and New Value columns show the original value and the last written value. If the write operation failed or the previous value could not be obtained, the corresponding field will display N/A. If the write operation failed, for example, if the device did not respond, you can select the item in the table and click Retry.

Click the Undo button to undo the selected successful attribute value change in results list if necessary.

User-Defined Attributes

In the Attribute Editor dialog, you can create a list of attributes that display when you expand the User Defined category. After you create this list, you can access the user-defined attributes. You can at any time remove the user defined attributes.

Create User-Defined Attributes

Each OneClick user can create a unique set of user-defined attributes. You can select user-defined attributes using the Attribute Selector dialog.

Follow these steps:

1. In the left panel of the Attribute Editor dialog, next to the User Defined folder, click Add.

The Attribute Selector dialog opens.

2. If you have more than one model type selected, select the model type whose attributes you want to edit from the left pane of the Attribute Selector dialog.

The attributes for the selected model type appear in the right pane of the Attribute Selector dialog.

3. Select the attribute to edit from the list, and click OK.

Note: Use the Filter text box to quickly locate an attribute or model type in the list.

The attribute that you selected appears in the User Defined category in the Attribute Editor dialog. You can only add attribute at a time to the User Defined category.

4. Repeat this process to select additional user-defined attributes.

Note: Remove user-defined attributes by clicking the remove link next to the attribute that you want to remove.

More information:

[Entity Table Interface Stacking](#) (see page 125)

[Edit Attributes for Specific Devices or for Model Types](#) (see page 180)

Change Attributes in Conjunction with Search

You can use the Attribute Editor feature in conjunction with the Search feature in the Locator tab (Locator Search). By using the Locator Search feature with the Attribute Editor feature you can locate all models meeting certain criteria and attempt to change the attribute values on those matching models.

The following example combines ‘creating and running a new search’ with ‘changing attributes through the Attribute Editor.’ It walks you through adding user-defined attributes and writing changes to the component (that is, SpectroSERVER, devices, interfaces) meeting the search criteria.

Note: For more information about using the Search feature in the Locator tab, see the *Operator Guide*.

Example: Define a Search to Create an Attribute for Editing

This example demonstrates how to create and run a search to locate the GlobalConfig model on the SpectroSERVER. It then shows how to use the Attribute Editor to add the HibernationCommSuccessTries attribute to the User Defined category so that you can update the value as needed.

Note: The value for HibernationCommSuccessTries determines the number of successful attempts the SpectroSERVER must make to devices in hibernation mode before the devices can resume normal management communication. By default, the value of this attribute is 3.

To define a search to create a user-defined attribute for edit

1. In the OneClick Locator tab, click  (Create a new search).
The Create Search dialog opens.
2. Select ‘Model Type Name (0x10000)’ from the Attribute drop-down list.
3. Select ‘Equal To’ from the Comparison Type drop-down list.
4. Type **GlobalConfig** in the Attribute Value field.
5. Click Save As, type a name for the search (for example, ‘Hibernation attempts’), and click OK.
6. Click OK in the Create Search dialog.
7. Select the search you just created (‘Hibernation attempts’) in the Locator tab and click  (Launch the selected search).
The Select Landscape to Search dialog opens.
8. Select the landscapes to search and click OK.
The search results appear in the Results tab.
9. Right-click the GlobalConfig entry and select Utilities, Attribute Editor.
The Attribute Editor dialog opens.

10. Click the add link in the left panel of the Attribute Editor dialog, next to the User Defined folder.

The Attribute Selector dialog opens.

11. Click the folder named Other in the left panel of the Attribute Selector dialog.

12. In the Filter text box (below the left panel) type **GlobalConfig** and select the GlobalConfig entry under the 'Other' folder.

13. In the Filter text box under the right panel of the Attribute Selector dialog, type **HibernationCommSuccessTries**.

The HibernationCommSuccessTries attribute appears in the Attribute for GlobalConfig list in right panel.

14. Double-click the HibernationCommSuccessTries entry in the list to add it to the User Defined category.

15. In the Attribute Editor dialog, edit the user-defined attribute value by selecting it in the left panel and clicking the right arrow button to move its associated attribute fields to the editing panel.

16. In the right panel, edit the attribute values as desired then click Apply to write the changes to the component.

The Attribute Edit Results dialog opens listing the results of the changes made.

Edit Attributes for Specific Devices or for Model Types

This section provides examples for changing an attribute value for a specific model type, or for a specific set of devices.

Example 1: Edit Interface_Polling_Interval for Cisco Devices Supporting IPsec

Cisco IPsec tunnel interface management is available in CA Spectrum for Cisco devices that support the IPsec related MIBs. Once modeled, the tunnel models are updated every hour. If your environment requires less or more frequent updates to the tunnel models, you can use the Attribute Editor to do so.

The attribute Interface_Polling_Interval defines how frequently CA Spectrum monitors the MIB associated with the tunnel interface models so that the modeling is up to date. To turn off this monitoring, you must set the Interface_Polling_Interval attribute to 0. To change the frequency at which CA Spectrum monitors these MIBs, change the value for the attribute to the desired number of seconds between polling cycles. 3600 seconds equals one hour.

Edit the Interface_Polling_Interval attribute for Cisco devices.

Follow these steps:

1. Locate the Cisco routers on your network by creating a new Global Collection of Cisco routers that have the Interface_Polling_Interval set to 3600 seconds.
2. Select the Cisco routers on which you want to edit the Interface_Polling_Interval value from the global collection.
3. Right-click and select Utilities, Attribute Editor.
4. Add the Interface_Polling_Interval to the User Defined attributes list as described in [Create User-Defined Attributes](#) (see page 178).
5. Move the Interface_Polling_Interval attribute into the right-hand panel for editing.
6. Enter the value in seconds to which you want to change the polling interval.
7. Select OK.

The Attribute Edit Results dialog opens, displaying the results for each device you attempted to change the attribute value on.

8. If the change failed on any of the selected devices, select them and click Retry.
9. Close the Attribute Edit Results dialog, and click OK to close the Attribute Editor dialog.

Example 2: Edit the DeviceTypeDiscEnable Attribute for Specific Devices

The DeviceTypeDiscEnable attribute is used to allow or prevent changes to the device type name value for device models and model types. You can modify the value for this attribute.

To edit DeviceTypeDiscEnable for specific models, you first must add the DeviceTypeDiscEnable attribute to the User Defined category. Once you have done that you can begin to select the devices on which you want to prevent device type name customizations, as described in the following procedure.

Edit DeviceTypeDiscEnable for specific device models.

Follow these steps:

1. Add the DeviceTypeDiscEnable attribute to the User Defined category using the procedure described in using the procedure described in [Create User-Defined Attributes](#) (see page 178).
2. Select the Locater tab in the Navigation panel.
3. Expand the Devices folder and double-click By Model Name.

The Search dialog opens.

4. Enter the name of the device type model on which you want to prevent device type name customizations; select all applicable landscapes as necessary and click OK.

The devices using the device type model specified appear in the Lists tab.

5. Select the specific devices on which you want to prevent customizations.
6. Right-click and select Utilities, Attribute Editor to launch the Attribute Editor.
7. Expand the User Defined folder, select DeviceTypeDiscEnable, and click the right arrow to place the attribute in the right-side editing panel.
8. Set the attribute value to No and click Apply.
9. Verify that 'Set as Default' is not selected.

The Attribute Edit Results dialog opens and displays the results of the edit, either successful or unsuccessful. If successful, the DeviceTypeDiscEnable attribute is now set to false, or no, on the devices you selected in Step 5.

Example 3: Edit DeviceTypeDiscEnable for a Model Type

You can set an attribute value and apply it to the model type in the CA Spectrum modeling catalog and to all device models. Also, all device models created in the future using the model type will use the attribute value that you set in this manner.

Edit DeviceTypeDiscEnable for a model type.

Follow these steps:

1. Add the DeviceTypeDiscEnable attribute to the User Defined category using the procedure described in [Create User-Defined Attributes](#) (see page 178).
2. Create a search using the Locator tab in the Explorer that finds some or all the device models that use the new custom model type.
3. Select one of the device models using the new custom model type in the search results list, right-click it and select Utilities, Attribute Editor.

The Attribute Editor opens with the context of the select device.

4. In the Attribute Editor, select the DeviceTypeDiscEnable attribute in the User Defined category and move it to the Attribute Editing panel by clicking the right arrow.
5. Set the attribute value to No.

6. Select Set as Default to apply this change to the CA Spectrum catalog for the model type used by the device model selected in Step 3.

7. Click Apply.

A warning message appears.

8. Click Yes.

This action sets the attribute value for all the device models using the model type, and applies it to the model type in the CA Spectrum modeling catalog and to all device models. All future device models created using this model type will have the DeviceTypeDiscEnable attribute set to No, and the Device Type Attribute cannot be overwritten.

More information:

[Calculate Normalized CPU Utilization](#) (see page 197)

[Calculate Normalized Memory Utilization](#) (see page 198)

Model Type Reevaluation

In cases where devices have been replaced on your network, IP addresses may be assigned to new devices without your knowledge. Therefore, device models periodically verify that they are modeled using the correct model type. If the model type no longer matches the device identify, an alarm is generated on the model.

By default, this reevaluation of model types occurs every 24 hours. You can change this setting for all models.

Edit the Model Type Reevaluation Interval

You can modify the interval for model type reevaluation. The MTypeVerifyInterval attribute on the VNM model determines the reevaluation interval, which is every 24 hours by default. Set this value to 0 to disable model type reevaluation.

Follow these steps:

1. Select the VNM model and click the Attributes tab in the Component Detail panel.

The available attributes for the VNM model appear on the left side of the panel.

2. Locate the MTypeVerifyInterval attribute by typing it in the Filter field.

The MTypeVerifyInterval attribute appears on the left side of the panel.

3. Double-click the MTypeVerifyInterval attribute.

The MTypeVerifyInterval attribute and its value appear on the right side of the panel.

4. Double-click the MTypeVerifyInterval attribute on the right side of the panel.
The Edit dialog opens.
5. Clear the 'No Change' check box in the Edit dialog and type **0** in the field.
The attribute value is changed to 0.
The Attribute Edit Results dialog indicates whether the modification succeeded.
6. Click Close.
The Attribute Edit Results dialog closes; model type reevaluation has been disabled for all models.

Change Management Attributes

This group of attributes lets CA Spectrum maintain up-to-date configuration information about modeled devices. The following attributes in this grouping let CA Spectrum interrogate a device and gather information about its interfaces and connections after a specific event occurs:

- Automatically Reconfigure Interfaces
- Discovery After Reconfigure
- Discover Connections after Link Up Events
- Topologically Relocate Model

See [Update Device Interface and Connection Information](#) (see page 119) for more information about these attributes, and about configuring CA Spectrum to maintain updated device configurations.

The following attributes in this grouping configure CA Spectrum to have redundant ways of contacting devices:

- Enable Redundancy
- Generate Redundancy Alarms

CA Spectrum will use multiple IP addresses in a cascading manner to contact devices if a device fails to respond to queries made on its primary address. The devices must have multiple IP addresses configured in their IP tables.

More information:

[Redundant Connections Between CA Spectrum and Modeled Devices](#) (see page 128)
[Attribute Editor](#) (see page 127)

Interface Configuration Attributes

You set the value for a set of model type interface configuration attributes in the Interface Configuration grouping. Some of the attributes available include the following:

Admin Status

This attribute sets the administrative status for an interface.

Create Sub-Interfaces

Set this attribute to Yes, and if a device supports RFC1573, CA Spectrum will model the device's sub-interfaces.

Interface Name Primary Suffix

CA Spectrum uses the value of this required attribute to name interfaces for the model or models for the Rename Interface Models action. Choose from a set of available suffixes in the drop-down list when editing this attribute.

Interface Name Secondary Suffix

CA Spectrum uses the value of this optional secondary attribute to name interfaces for the model or models for the Rename Interface Models action. It is prefixed by an underscore (_) and follows the value of the primary suffix. This secondary attribute is optional.

More information:

[Rename Interface Models](#) (see page 125)

[Attribute Editor](#) (see page 127)

Stale Interfaces

CA Spectrum handles interface definitions that are temporarily removed from their corresponding MIB tables using its Stale Interface functionality. These types of situations can include:

- A module is temporarily removed from a device
- A configured tunnel temporarily goes down on a device

In these situations, it is advantageous for CA Spectrum to retain the interface modeling information rather than immediately destroying them. This prevents useful model-specific attributes and resolved connections from being lost.

CA Spectrum determines that an interface model is stale when no corresponding entry exists in the MIB where the interface is defined. Once the stale interface 'ages out' CA Spectrum removes it from the model. The age out period for an interface is defined by the Stale Interface Age Out attribute (in minutes) on the device model.

An event is generated when CA Spectrum determines an interface is stale. If the stale interface has resolved connections, a minor alarm is generated on the interface model. If CA Spectrum determines that the interface is no longer stale before the age out period expires, and prior to a reconfiguration which causes the interface model to be destroyed, an event is generated and any stale interface alarm is cleared.

Enable Stale Interface Alarms

Set this attribute to Yes to enable CA Spectrum to generate a minor alarm on an interface model when the interface becomes stale and it has resolved connections. Set this attribute to No if you do not want this condition to generate an alarm.

Stale Interface Age Out (min)

This attribute specifies the amount of time in minutes that CA Spectrum waits for the stale interface to 'age out' before removing the model. To disable the stale interface functionality, set this attribute to 0.

Default: 120

Maintenance Mode Attributes

You can set values for the Maintenance and Hibernation mode attributes in the Attribute Editor's Maintenance Mode folder. You can also set these attribute values and create and apply maintenance mode schedules in the model's Information tab.

Note: For more information about Maintenance and Hibernation modes and managing maintenance mode schedules, see the *Operator User Guide*.

Rollup Alarm Attributes

Access the CA Spectrum attributes used to manage rollup alarm settings (conditions) and threshold levels in the Attribute Editor's Roll-Up Alarm Settings folder. You can also view and set these attributes in the Information tab under the General Information subview. You can adjust these attributes for containers modeled on your network and for the CA Spectrum container model library. The following section lists the attributes, describes how they are used, and defines their default values.

Note: Change threshold levels carefully; you may see an increase in generated alarms if threshold levels are set lower, or a decrease in generated alarms if levels are set higher.

Value When Yellow

The point value of a Yellow alarm condition existing in a child towards the roll up alarm threshold value for the parent container.

Default: 1

Value When Orange

The point weight of an Orange alarm condition existing in a child towards the roll up alarm threshold value for the parent container.

Default: 3

Value When Red

The point weight of a Red alarm condition existing in a child towards the roll up alarm threshold value for the parent container.

Default: 7

Yellow Threshold

The minimum points needed to trigger a Yellow roll up alarm for a container.

Default: 3

Orange Threshold

The minimum points needed to trigger an Orange roll up alarm for a container.

Default: 6

Red Threshold

The minimum points needed to trigger a Red roll up alarm for a container.

Default: 10

More information:

[General Information Subview](#) (see page 149)

[Rollup Condition Thresholds](#) (see page 188)

Model Status and Alarm Conditions

OneClick uses rollup alarm thresholds and model alarm thresholds to determine the status for modeled entities. OneClick displays two types of status for modeled entities, Condition and Rollup Condition. The following section lists the details about what these conditions describe, and what OneClick applies them to.

Condition

Applies to all models. Reflects the current contact or alarm status of the model itself.

Rollup Condition

Applies to container models, such as networks, LANs, and WANs. Reflects the composite status of all the other models in the container, which are sometimes referred to as their children.

If a modeled device or interface exists in a container, its condition rolls up to the parent container and is reflected in the container's rollup condition. The model status types listed previously rely on threshold values to determine when and how to use the associated color indicators. The rollup condition is displayed using an inverted triangle that appears adjacent to a container icon in the container's Information tab.


Rollup Condition Thresholds

A container model has attributes that define values for alarm conditions that may exist on the children of the container. A container model also has attributes that define when its rollup alarm conditions are triggered. The combined value of all the alarm conditions for a container's children is used to determine the rollup alarm condition for the container.

The following illustration shows a container that has a rollup condition of orange, or major, based on the alarm conditions of its two children. The rollup alarm setting for the container uses the default values listed previously.

- One child has a green or normal condition; this contributes zero points toward the container's rollup condition.
- The other child has a red or critical condition that contributes seven points toward the container's rollup condition (Value When Red = 7).

- The total value of the alarm conditions on the container's children is 7.
- The rollup thresholds for the container use the default values for the Rollup alarm settings listed in [Rollup Alarm Attributes](#) (see page 186). The Orange Threshold value = 6, so the container's rollup appears as Orange, indicating a major alarm condition.



SNMP Communication Attributes

You can tune overall SNMP communications by changing the values of the attributes in the Attribute Editor's SNMP Communication folder. The following attributes define how CA Spectrum communicates with a device:

SNMP Community String

Lets the SpectroSERVER communicate with devices on your network.

DCM Timeout (ms)

The number of milliseconds the polling agent will wait for a response from the device before timing out.

DCM Retry Count

Specifies the number of times the SpectroSERVER retries to establish device communication after the DCM timeout value expires.

Polling Interval (sec)

The number of seconds between polls CA Spectrum makes to devices.

Note: Increasing this number results in less SNMP-related traffic on your network and a smaller load on the SpectroSERVER. Decreasing this number for mission critical devices and interfaces lets you see updated information about these devices in OneClick more often. This can improve your ability to see potential issues on the network before they affect network performance. A decreased Polling Interval will result in more SNMP network traffic generated by CA Spectrum.

Poll To Log Ratio

The number of polls per log. If it is set to 3, then data is logged every third poll.

More information:

[Discover Connections After Link-Up Events](#) (see page 120)

Threshold Attributes

The Thresholds grouping contains the CA Spectrum device and interface threshold settings.

Note: For more information about interface threshold parameters, see the *Operator Guide*.

More information:

[Device Threshold Settings](#) (see page 115)

How CA Spectrum Calculates CPU and Memory Utilization

The MIB objects that are read to calculate the source of the CPU and memory utilization can be customized for an individual device model or set as the default for a given model type. This is only necessary if CA Spectrum is unable to identify a source for CPU and memory utilization, or if you prefer to use a different source.

Do one of the following to customize CPU and memory utilization:

- Modify the order and the types of sources to be tested.
- Modify the source attributes used by the attribute redirection type.

Attribute redirection is the process of using well-known attributes as a pointer, or redirectors, to proprietary attributes. For example, the well-known attribute `NRM_CPUUtilAttr (0x12e2d)` on the `AirespaceSw` model type holds the attribute ID of the Airespace proprietary CPU attribute, `agentCurrentCPUUtilization (0x4b605ae)`.

With attribute redirection, it is possible for generic code to reference unique attributes for each device model or model type. Additionally, attribute redirection lets you change the source attributes without having to restart the SpectroSERVER. By default, for most model types, attribute redirection is the first source type that is tested. Therefore, in most circumstances, you only have to modify the attributes that attribute redirection is using.

Note: Only certain attribute types can be used for attribute redirection.

Before you make changes to the normalized CPU and memory intelligence, understand how CA Spectrum calculates CPU and memory utilization.

CA Spectrum does the following to calculate CPU and memory utilization:

1. CA Spectrum identifies the source that is used to calculate the CPU and memory utilization. CA Spectrum identifies the source any time the device is reconfigured.

A list of the possible sources is provided in a new preference attribute. Each source is tested in order. If a valid source is found, the source intelligence, the attributes that are used, and the attribute's model type are stored for reuse during the utilization calculation. This helps ensure that the source does not change until the source is reconfigured. If the list of sources is empty, or a valid, functioning source is not identified, the source type is set to "None" and no further reads are made to the device.

In general, the sources are tested in the following order:

- Attribute redirection
- CA proprietary intelligence
- Standard intelligence (RFC 2790 and Net-SNMP)

2. CA Spectrum performs the actual calculation of the utilization using the correct source, attribute IDs, and model type handles that were identified. The running attribute IDs list and the running model handle list are passed into the calculation method each time, to help ensure that the same attributes are read.

More information:

[Normalized Memory Utilization Attributes](#) (see page 195)

[Normalized CPU Utilization Attributes](#) (see page 193)

[Device Threshold Settings](#) (see page 115)

Normalized CPU Utilization Calculation Requirements

You can calculate the normalized CPU utilization for any device using attribute redirection as the source for the utilization, however, the device must meet the following requirements:

- The device must have a single MIB object which has a data type of either integer, 64-bit integer, text string, float, or real. The MIB object must be a scalar object or a list object.
- If the MIB object has a data type of text string, the values of the MIB object are valid if the text string represents a valid number. For example, 9.4, 43, and 1200, are considered valid text strings. A text string that contains numbers with extra text is not valid. For example, 43% is considered invalid.
- If the MIB object is a list object, each instance in the list must report a valid CPU value (no filtering of lists is provided). The reported value must be an instantaneous usage (or an aggregate of a short time period).
- The MIB object must report the utilization of all CPUs in the device in units of 0-100 percent.

Note: CA Spectrum does not attempt to adjust invalid results. For example, if an attribute returns 110 percent utilization, CA Spectrum reports the 110 percent utilization. Verify that the values that are reported in the attribute always yield the correct utilization, once calculated. If negative values are returned, thresholds can indicate a violation, even though a threshold was not exceeded.

More information:

[Calculate Normalized CPU Utilization](#) (see page 197)

[Device Threshold Settings](#) (see page 115)

Normalized Memory Utilization Calculation Requirements

You can calculate the normalized memory utilization for any device, using attribute redirection as the source for the utilization, however, the device must meet the following requirements.

The device must have one of the following:

- A single MIB object which is of the data type integer, 64-bit integer, text string, float, or real, and is either a list or a scalar. In addition, this object must report the utilization of all memory for the device in units of 0-100 percent.

Note: CA Spectrum does not attempt to adjust invalid results. For example, if an attribute returns 110 percent utilization, CA Spectrum reports the 110 percent utilization. Verify that the values that are reported in the attribute always yield the correct utilization, once calculated. If negative values are returned, thresholds can indicate a violation, even though a threshold was not exceeded.
- If the MIB object has a data type of text string, the values of the MIB object are valid if the text string represents a valid number. For example, 9.4, 43, and 1200, are considered valid text strings. A text string that contains numbers with extra text is not valid. For example, 9.4 MB, 43 MB, and 1,200 are considered invalid.
- Two or more MIB objects which are all either scalars or lists. Additionally, two of the objects must report either the total amount of free memory, used memory, or total memory. The units in which each MIB object reports its respective value *must* be the same.

Note: CA Spectrum does not attempt to verify that the units in which each MIB object reports its respective value are the same.

The values reported by the MIB objects must be instantaneous (or an aggregate of a short time period).

More information:

[Calculate Normalized Memory Utilization](#) (see page 198)

[Device Threshold Settings](#) (see page 115)

Normalized CPU Utilization Attributes

CA Spectrum uses the following attributes to calculate the normalized CPU utilization for a device:

NRM_CPUIntelPref

Lists possible sources to test when identifying normalized CPU utilization. These sources are tested in the order in which they appear in this attribute.

NRM_DeviceCPUUtilization

Reports the device's CPU utilization. The normalized CPU utilization calculation is triggered based on what this attribute reports.

NRM_DeviceCPUUtilizationNames

Contains the names of each instance of the CPU utilization value. By default, the instance is displayed as:

CPU: <instance>

<instance>

Is the instance ID of each CPU utilization value.

Note: If the NRM_CPUUtilNameAttr attribute is available, the NRM_CPUUtilNameAttr attribute setting is used to populate the NRM_DeviceCPUUtilizationNames attribute. The names of each instance of the CPU utilization value is set once. If the NRM_CPUUtilizationNameAttr attribute setting changes, you must reconfigure the model to pick up the name change.

NRM_CPUAttr_Source

Contains the source that is currently being used to calculate the normalized CPU utilization. If the intelligence ID is attribute redirection, this attribute says 'Attribute Redirection'. If the intelligence ID is CA – Proprietary, this attribute lists the MIB that CA Spectrum reads the values from.

The following attributes are used to utilize attribute redirection as the source for calculating the normalized memory utilization for a device:

NRM_CPUUtilAttr

(Required) Points to the attribute that reports CPU utilization in percent. You populate this attribute with the attribute ID of the attribute which reports CPU utilization for the device. The attribute can be either a list or a scalar, and must be one of the following data types: counter, gauge, int, real, or 64-bit long.

Note: CA Spectrum does not attempt to adjust invalid results. For example, if an attribute returns 110 percent utilization, CA Spectrum reports the 110 percent utilization. Verify that the values that are reported in the attribute always yield the correct utilization, once calculated. If negative values are returned, thresholds can indicate a violation, even though a threshold was not exceeded.

NRM_CPUUtilNameAttr

(Optional) Points to the attribute that reports the identifying information for the CPUs. This attribute holds the attribute ID of the attribute which reports the names associated with each instance of CPU utilization for this device.

This attribute can be a scalar or a list, but it must be the same data type as the NRM_CPUUtilAttr attribute and the names must be ordered such that the first element in the list matches the first element in the utilization list. If you do not enter an attribute ID, or if the attribute ID you provide is not valid, the instance ID is appended to "CPU:<name>" to create the name for that CPU value. Any data type is accepted.

Note: OctetStrings are treated as printable text strings.

NRM_CPUModelTypeToRead

(Optional) Lists the model type handle of an application model that the NRM_CPUUtilAttr and the NRM_PUUUtilNameAttr attributes are read from.

More information:

[How CA Spectrum Calculates CPU and Memory Utilization](#) (see page 191)

[Troubleshoot CPU and Memory Utilization Calculation](#) (see page 200)

[Device Threshold Settings](#) (see page 115)

Normalized Memory Utilization Attributes

CA Spectrum uses the following attributes to calculate the normalized memory utilization for a device:

NRM_MemoryIntelPref

Lists possible sources to test when identifying normalized memory utilization. These sources are tested in the order in which they appear in this attribute.

NRM_DeviceMemoryUtilization

Reports the device's memory utilization. The normalized memory utilization calculation is triggered based on what this attribute reports.

NRM_DeviceMemoryUtilizationNames

Contains the names of each instance of the memory utilization value. By default, the default is displayed as:

Memory: <instance>

<instance>

Is the instance ID of each memory utilization value.

Note: If the NRM_MemoryUtilNameAttr attribute is available, the NRM_MemoryUtilNameAttr attribute setting is used to populate the NRM_DeviceMemoryUtilizationNames attribute. The names of each instance of the memory utilization value is set once. If the NRM_MemoryUtilizationNameAttr attribute setting changes, you must reconfigure the model to pick up the name change.

NRM_MemAttr_Source

Contains the source that is currently being used to calculate the normalized memory utilization. If the intelligence ID is attribute redirection, this attribute says 'Attribute Redirection'. If the intelligence ID is CA – Proprietary, this attribute lists the MIB that CA Spectrum reads the values from.

The following attributes are used to utilize attribute redirection as the source for calculating the normalized memory utilization for a device:

NRM_MemoryUtilAttr

Points to the attribute that reports the memory utilization in percent.

NRM_MemoryUsedAttr

Points to the attribute that reports the used memory in units, for example, bytes, kilobytes, megabytes, gigabytes, and so on.

NRM_MemoryTotalAttr

Points to the attribute that reports the total memory in units, for example, bytes, kilobytes, megabytes, gigabytes, and so on.

NRM_MemoryFreeAttr

Points to the attribute that reports the free memory in units, for example, bytes, kilobytes, megabytes, gigabytes, and so on.

Note: To calculate the normalized memory utilization for a device using attribute redirection as the source for the utilization, either NRM_MemorUtilAttr must be populated, or two of the following three attributes must be populated: NRM_MemoryUsedAttr, NRM_MemoryTotalAttr, NRM_MemoryFreeAttr.

NRM_MemoryUtilNameAttr

(Optional) Points to the attribute that provides the memory utilization names.

NRM_MemoryModelTypeToRead

(Optional) Lists the model type handle of an application model that the NRM_MemoryUtilAttr, NRM_MemoryUsedAttr, NRM_MemoryTotalAttr, NRM_MemoryFreeAttr, and the NRM_MemoryUtilNameAttr attributes should be read from.

More information:

[How CA Spectrum Calculates CPU and Memory Utilization](#) (see page 191)

[Troubleshoot CPU and Memory Utilization Calculation](#) (see page 200)

[Device Threshold Settings](#) (see page 115)

Calculate Normalized CPU Utilization

You can calculate the normalized CPU utilization for a device whose utilization is not calculated out-of-the-box. You can also recalculate the utilization using different attributes than those attributes that are used by default. You can use different attributes on a per model or a per model type basis. To use different attributes on a per model type basis, use the Attribute Editor or the Model Type Editor to change the default attribute values for the model types.

Note: For more information about the Model Type Editor, see the *Model Type Editor User Guide*.

To determine if the normalized CPU utilization is already calculated for a device, view the Source column in the Thresholds and Watches, Thresholds subview in the Information tab of the Component Details panel.

Note: CA Spectrum does not attempt to adjust invalid results. For example, if an attribute returns 110 percent utilization, CA Spectrum reports the 110 percent utilization. Verify that the values that are reported in the attribute always yield the correct utilization, once calculated. If negative values are returned, thresholds can indicate a violation, even though a threshold was not exceeded.

Follow these steps:

1. Verify that the device meets the specified [requirements](#) (see page 192).
2. Identify the attribute that reports CPU utilization and place this attribute ID into NRM_DeviceCPUUtilAttr.
3. Identify the attribute that reports the identifying information for all the device CPUs, if the attribute exists.
4. Place the attribute ID into NRM_CPUUtilNameAttr.
5. Identify the origin of the attribute that reports CPU utilization and the attribute that reports the identifying information for the CPUs. If these attributes originate on an application model, enter the model type handle of the application model into NRM_DeviceCPUModelTypeToReadAttr. Otherwise, leave this attribute empty.

If the DeviceCPUModelTypeToReadAttr is empty, CA Spectrum attempts to read the specified attribute from the device model. If DeviceCPUModelTypeToReadAttr is populated, CA Spectrum attempts to find an associated application model with that model type handle. If the associated application model, with the specified model type handle, is not found, or the attributes do not exist on that model, attribute redirection is not considered to be a valid source for calculating the utilization.

6. Reconfigure the device model.

If attribute redirection fails, CA Spectrum attempts to test other available sources. If a valid, functioning source is not identified, the source column in the Thresholds and Watches, Thresholds subview displays 'None'. The normalized CPU and normalized memory performance graphs report 'Not available', and no further reads are made to the device. If a valid, functioning source is identified, the source column displays the successful source.

More information:

[Edit Attributes for Specific Devices or for Model Types](#) (see page 180)
[Device Threshold Settings](#) (see page 115)

Calculate Normalized Memory Utilization

You can calculate the normalized memory utilization for a device whose utilization is not calculated by default, or for a device whose utilization has not been automatically calculated. You can also recalculate memory utilization using attributes that are not used by default. You can use different attributes for individual models or for selected model types. To use different attributes on a per model type basis, use the Attribute Editor or the Model Type Editor to change the default attribute values for the model types.

Note: For more information about the Model Type Editor, see the *Model Type Editor User Guide*.

To determine whether the normalized memory utilization is already calculated for a device, view the Source column in the Thresholds and Watches, Thresholds subview in the Information tab of the Component Details panel.

Note: CA Spectrum does not attempt to adjust invalid results. For example, if an attribute returns 110 percent utilization, CA Spectrum reports the 110 percent utilization. Verify that the values that are reported in the attribute always yield the correct utilization, once calculated. If negative values are returned, thresholds can indicate a violation, even though a threshold was not exceeded.

Follow these steps:

1. Verify that the device meets the specified [requirements](#) (see page 193).

2. Identify the attribute that reports memory utilization, if it exists, and place this attribute's handle into `NRM_DeviceMemoryUtilizationAttr`. If the device does not support an attribute that reports memory utilization, identify two of the following three attributes:

- Used Memory
- Free Memory
- Total Memory

Populate the related attribute, `NRM_MemoryXXXAttr`, where `XXX` is Free, Used, or Total. If more than the required attributes are provided, all the attributes must exist on the same model type.

Note: These attributes must report memory utilization in the same unit, for example, bytes, kilobytes, megabytes, gigabytes. However, CA Spectrum does not verify that each attribute reports the same units.

If more than the required number of attributes are provided, CA Spectrum uses the following order of precedence to determine which attributes to use:

1. `NRM_MemoryUtiliAttr`
2. `NRM_MemoryUsedAttr` and `NRM_MemoryTotalAttr`
3. `NRM_MemoryFreeAttr` and `NRM_MemoryTotalAttr`
4. `NRM_MemoryFreeAttr` and `NRM_MemoryUsedAttr`

The first instance that returns a valid set of values is used.

3. Identify where these attributes originate. If these attributes originate on an application model, enter the model type handle of the application model into `NRM_DeviceMemoryUtilizationNameAttr`. Otherwise, leave this attribute empty.

If the `NRM_DeviceMemoryUtilizationNameAttr` is left empty, CA Spectrum attempts to read the specified attribute from the device model. If `NRM_DeviceMemoryUtilizationNameAttr` is populated, CA Spectrum attempts to find an associated application model with that model type handle. If the associated application model, with the specified model type handle, is not found, or the attributes do not exist on that model, attribute redirection is not considered to be a valid source for calculating the utilization.

4. Reconfigure the device model.

If attribute redirection fails, CA Spectrum attempts to test other available sources. If a valid, functioning source is not identified, the source column in the Thresholds and Watches, Thresholds subview displays 'None', the performance graphs report 'Not available', and no further reads are made to the device. If a valid, functioning source is identified, the source column displays the successful source.

More information:

[Edit Attributes for Specific Devices or for Model Types](#) (see page 180)
[Device Threshold Settings](#) (see page 115)

Troubleshoot CPU and Memory Utilization Calculation

If CA Spectrum is returning incorrect or invalid CPU and memory utilization calculation values, reconfigure the model.

The attributes that CA Spectrum reads, and the model handles that CA Spectrum reads from, are all cached to verify that the same source is used during the utilization calculation. If a source is no longer available, or has been changed, invalid or incorrect values are reported until the device model is reconfigured.

Note: Name attributes are only reevaluated when the model is reconfigured and then cached. Therefore, if instances change, or the names associated with a given instance change, reconfigure the model.

If CA Spectrum is not selecting the source that you want to use, do the following:

1. Verify that the order of possible sources you want to test is set correctly in the CPU intelligence preference attribute, [NRM_CPUIntelPref](#) (see page 193), and memory intelligence preference attribute, [NRM_MemoryIntelPref](#) (see page 195).
2. Verify that the device supports the source you want to use. For example, if you want to use attribute redirection, confirm that the attributes are supported on the device model or specified application model.

Chapter 7: Fault Management

This section contains the following topics:

[Fault Isolation Settings](#) (see page 201)

[Port Fault Correlation](#) (see page 203)

[Configure Cross-Landscape Fault Correlation](#) (see page 212)

[Wide Area Link Monitoring](#) (see page 221)

[Port Layer Alarm Suppression](#) (see page 224)

[Port Criticality](#) (see page 224)

[Live Pipes and Fault Management](#) (see page 224)

[Suggested Port Fault Settings for Optimal Fault Notification](#) (see page 229)

[Device Criticality](#) (see page 230)

[Configuring Fault Management for Pingables](#) (see page 230)

[False Management Lost or Contact Lost Alarms](#) (see page 233)

Fault Isolation Settings

The Fault Isolation subview in the VNM model's Information tab lets you configure various aspects of the CA Spectrum device fault isolation functionality. It contains the following settings:

ICMP Support Enabled

Specifies whether an attempt should be made to contact a device using the ICMP protocol when trying to ascertain the fault status of the device. When the ICMP_SUPPORT attribute is enabled (set to TRUE), CA Spectrum looks at the setting of the ICMP_SUPPORT attribute at the device level. If ICMP support is also enabled on the device, CA Spectrum attempts to contact the device using the ICMP protocol. However, when ICMP_SUPPORT is disabled (set to FALSE), this setting takes precedence over the setting at the device level and prevents any attempt to contact the device using the ICMP protocol for fault isolation.

Default: Yes

ICMP Timeout (msec)

Specifies the amount of time (in milliseconds) CA Spectrum waits for a response to an ICMP ping. If a response is not received within this period of time, CA Spectrum assumes the device has timed out.

Default: 3000 milliseconds (3 seconds)

ICMP Try Count

Specifies the total number of attempts made to contact a device using the ICMP protocol before CA Spectrum determines the device is down.

Lost Device Try Count

Specifies the number of retries that CA Spectrum attempts for each SNMP request sent to a device after contact with the device is lost.

Default: 1

Port Fault Correlation

Enables port fault correlation and specifies how it should be configured.

Default: All Connected Ports

Contact Lost Model Destruction

Specifies whether a device is automatically destroyed when its CONTACT_STATUS is set to false. When enabled, models that have had their CONTACT_STATUS set to lost for a specified period of time, as determined by the Destruction Delay (sec) setting, are destroyed automatically. When disabled, models will never be automatically destroyed as a result of the value of the CONTACT_STATUS attribute.

Default: Disabled

Destruction Delay (sec)

Specifies the length of time (in seconds) that a model must continuously have its CONTACT_STATUS attribute set to lost before it is automatically destroyed.

Default: 604800 seconds (7 days)

Destruction Event Generation

This field controls whether an event message is created when a model is automatically destroyed. When this field is set to Enabled, an event will be generated each time a model is destroyed as a result of its CONTACT_STATUS being continuously set to lost for the length of time specified in the Destruction Delay Time field. When this field is set to Disabled, no event message will be generated.

Default: Enabled

Router Redundancy Retry Count

Specifies the number of times CA Spectrum will attempt to contact a router's redundant IP addresses if contact is lost with its primary address. The polling interval setting determines the amount of time between each attempt.

Default: 2

Unresolved Fault Alarm Disposition

If information about the connectivity of your network model is incomplete, CA Spectrum may be unable to find the root cause of a network outage. In this case, the status of all devices affected by the outage is set to gray, and a red unresolved fault alarm is generated. This alarm indicates that CA Spectrum has lost contact with a group of devices, but was unable to pinpoint the cause.

All of the devices to which CA Spectrum has lost contact are listed in the impact scope of the alarm. The model name and other details of the lost devices also appear in the event that generated the alarm.

The Unresolved Fault Alarm Disposition field lets you control how the unresolved fault alarm is generated. When set to Fault Isolation Model, the alarm will be generated on the Fault Isolation model. When set to Device In Fault Domain, the alarm will be generated on one of the devices with which CA Spectrum has lost contact. When determining which device to generate the alarm on, CA Spectrum looks for the device with the highest criticality. If the highest criticality is shared by two or more devices, CA Spectrum generates the alarm on the first of these devices that it finds. If all devices have the same criticality, then CA Spectrum chooses the device with the lowest model handle.

WA Link Fault Isolation Mode

Specifies whether WA_Link models are considered neighbors for fault isolation purposes. Options are Normal and Transparent.

Default: Normal

More information:

[Port Fault Correlation](#) (see page 203)

[Fault Isolation Subview](#) (see page 161)

Port Fault Correlation

CA Spectrum lets you customize its fault isolation algorithm to resolve the root cause of a network outage to the port level. This is most desirable in cases where a single physical port, such as a Frame Relay interface, supports multiple logical connections to remote devices. If the physical port goes down, CA Spectrum can suppress the alarms on all downstream devices in favor of a single red alarm on the physical interface, thus significantly reducing the number of alarms which need attention. The impact severity and scope of the red alarm on the physical interface will contain all downstream devices, as well as the physical interface.

Port Fault Correlation Options

Use the Port Fault Correlation setting in the VNM model's Fault Isolation subview to configure port fault correlation.

Disabled

Disables port fault correlation. The root cause of a network outage will remain a red alarm on a device model. However, Fault Isolation will still examine all of the device's connected ports to see if they are all in Maintenance Mode. If so, the alarm on the device model will be suppressed.

All Connected Ports

Port fault correlation will run, examining all ports that exist on "up" neighbors which are connected to the down device as possible root causes of the outage. No additional manual configuration is required.

Management Neighbors Only

Enables port fault correlation to run but only examine ports which were previously configured manually as management neighbors as possible root causes of the outage.

All Connected Ports—Multiple Devices Only

Enables port fault correlation, examining all ports that exist on "up" neighbors which are connected to the down device as possible root causes of the outage. However, CA Spectrum will only resolve the outage to the port level if there is more than one device model that would have a red alarm which can be correlated to the port alarm. If only one connected device alarm can be correlated to the port alarm, CA Spectrum will not suppress the device alarm. Instead, both the port and device alarm will be generated.

More information:

[Fault Isolation Subview](#) (see page 161)

[Fault Isolation Settings](#) (see page 201)

[Suggested Port Fault Settings for Optimal Fault Notification](#) (see page 229)

Port Fault Correlation Criteria

The following criteria must be met for the root cause of an outage to be resolved to the port level:

- The down device must have only one “up” neighbor. If the down device has more than one “up” neighbor, port fault correlation will not be performed. This is done to reduce the number of alarms for a single problem. If multiple up neighbors were a valid criteria, and all connected ports were down, multiple red alarms would exist, all with the same impact severity and scope. If a device has more than one up neighbor, CA Spectrum assumes the problem lies with the device, not the upstream ports and creates a single red alarm on the device.
- The down device must have at least one connected port (or management neighbor port) on an “up” neighbor that is down.
- If multiple ports on the “up” neighbor connect to the down device (such as link aggregation), all of the ports must be down.
- A port is considered “down” if it is either operationally down, or the port model has been put into Maintenance Mode.
- There must be an alarm on at least one of the down ports. Otherwise, there would be no alarm to which CA Spectrum could resolve the outage.
- If Port Fault Correlation is set to Management Neighbors Only, management neighbors for the down device must have been configured before the outage occurred.

Port Fault Correlation Caveats


Port Fault Correlation overrides the Suppress Linked Port Alarms setting in the Live Pipes subview. When set to TRUE, this setting suppresses the alarm on an upstream port if it's connected to an unreachable device. If Port Fault Correlation is enabled, and the upstream port is the root cause of an outage, CA Spectrum forces the upstream port to alarm.

Only the Criticality of the alarmed port will be used in the impact severity and scope calculation of the root cause alarm. The Criticality of any sub-interfaces (such as DLCI ports) will not be included.

Port Fault Correlation is supported by Device models only. Models such as Fanouts and Unplaced do not support this feature. WA_Link models have their own mechanism for supporting port fault correlation, Link Fault Disposition, which is explained in [Wide Area Link Monitoring](#) (see page 221).

If multiple ports on the “up” neighbor connect to the down device (e.g. link aggregation), and all of the ports are down, multiple red alarms will exist as the root causes of the outage. Each red alarm will contain the same impact severity and scope. The root cause of the outage in this case is all of the ports, not just one of them.


Example: Port Fault Correlation Scenario 1



The previous diagram assumes that CA Spectrum must communicate through Router A to reach Routers 1 through N, and that this is the only means by which CA Spectrum can reach them. Each remote router is connected to Router 1 using a frame relay link. In CA Spectrum, this is modeled by connecting each DLCI port model to the other device.

If the physical frame relay interface (FR A) goes down in this scenario, all virtual circuits provisioned on the interface will go down as well. With Port Fault Correlation disabled, the alarms shown in the following diagram will occur.


Fault Scenario 1: Alarms without Port Fault Correlation



If a trap is received for FR A going down (or a live pipe is configured to be on), the physical frame relay interface will have a red alarm on it. In addition, all routers connected to the frame relay interface will have a red alarm on them. This could mean multiple red alarms could be generated by CA Spectrum for a single problem.

Port Fault Correlation reduces the number of alarms generated for this problem to a single alarm without requiring any manual configuration beforehand. The following diagram shows the results with Port Fault Correlation enabled.

Fault Scenario 1: Alarms with Port Fault Correlation




A single red “Bad Link” alarm will be seen in the Alarms tab. That alarm will have an Impact Scope and Severity which contains the following models: FR A, Routers 1 through N, and all unreachable devices that are downstream from Routers 1 through N.

Example: Port Fault Correlation Scenario 2

This fault scenario illustrates the benefits of setting the Suppress Linked Port Alarms and Port Fault Correlation attributes as recommended in [Suggested Port Fault Settings for Optimal Fault Notification](#) (see page 229).


The following diagram assumes that the VNM must communicate through Routers A and B to reach Routers C, D, and E, and that is the only means by which the VNM can reach them. In CA Spectrum, port-level connectivity is modeled as shown.

Fault Scenario 2: Multiple “Up” Neighbors



In this scenario, Router C goes down, which causes CA Spectrum to lose contact with Routers C, D, and E, and makes Ports A and B go down as well. If Suppress Linked Port Alarms is set to TRUE, and Port Fault Correlation is set to All Connected Ports, only a single red alarm on Router C will result, as shown in the following diagram:


Fault Scenario 2: Multiple “Up” Neighbors



The upstream ports (Ports A and B) have their alarms suppressed because Suppress Linked Port Alarms is set to TRUE. Even though Port Fault Correlation is enabled, Router C has multiple “up” neighbors, so the fault won't be resolved to the port level. When this occurs, CA Spectrum assumes the fault is with the device itself, not the connected ports.

If you set Suppress Linked Port Alarms to FALSE, and Port Fault Correlation is still set to All Connected Ports, Router C and the upstream ports will be alarmed (if the status of the ports is being polled, or CA Spectrum receives a LinkDown trap), as shown in the following diagram:


Fault Scenario 2: Multiple “Up” Neighbors



Once again, the fault wasn't resolved to the port level because Router C has multiple “up” neighbors. Since Suppress Linked Port Alarms is disabled, CA Spectrum will alarm the upstream ports.


If Router C had only one “up” neighbor, as shown in the following diagram, even if Suppress Linked Port Alarms were set to TRUE (assuming Port Fault Correlation is still set to All Connected Ports), CA Spectrum will resolve the fault down to the port level. Port Fault Correlation forces the upstream port to be alarmed, and the alarm on Router C is suppressed.

Fault Scenario 2: Single “Up” Neighbor




Example: Port Fault Correlation Scenario 3

This fault scenario demonstrates what happens when Port Fault Correlation is set to All Connected Ports--Multiple Devices Only. It assumes that CA Spectrum must communicate through Router A to reach Routers 1 through N, and that this is the only means by which CA Spectrum can reach them. Each remote router is connected to Router 1 using a frame relay link. This is modeled by connecting each DLCI port model to the other device, as shown in the following diagram:




Assume the physical frame relay interface (FR A) goes down. This means that all virtual circuits provisioned on the interface will go down as well. With Port Fault Correlation disabled, the alarms shown in the following diagram will occur:

Fault Scenario 3: Alarms without Port Fault Correlation



With Port Fault Correlation set to All Connected Ports--Multiple Devices Only, the alarms in the following diagram will occur because multiple devices can be correlated to the frame relay interface.

Fault Scenario 3: Port Fault Correlation Set To “All Connected Ports--Multiple Devices Only”



With Port Fault Correlation set to All Connected Ports--Multiple Devices Only, if there is only one router lost because of a down link, then the alarm on the remote router will not be suppressed.



Port Fault Correlation Anomalies

If a red alarm is generated on a port model as the root cause of an outage, you may then choose to put that port model into Maintenance Mode. If so, the red alarm would be replaced with a brown alarm. The brown alarm will still contain the same impact severity and scope (except the maintenance port will no longer contribute to the impact). If you then decide to take the port out of Maintenance Mode, the red alarm will reappear. It is possible, in this scenario, for the impact scope and severity of the red alarm to be lost.

Configure Cross-Landscape Fault Correlation

In a distributed SpectroSERVER (DSS) environment, a network administrator may need to model a router from a local landscape in a remote landscape for its connections to participate in fault isolation for the remote landscape. This *proxy* model doesn't need to participate in alarm generation in the remote landscape because it already does so in the local landscape where it is modeled "normally." It is in the local landscape that alarms for the router are tracked and trouble tickets are created.

In such a scenario, Cross-Landscape Fault Correlation can prevent multiple red alarms for the same outage. With the Enable Event Creation attribute set to FALSE in the proxy model's Fault Management subview, CA Spectrum suspends the creation of events for the model (and any component models such as boards or ports). This effectively disables alarms for the model, but unlike maintenance mode, SNMP communication with the proxy model continues, and so too does its participation in fault isolation.

Note: For more information about distributed network management, see the *Distributed SpectroSERVER Administrator Guide*.

Designate a Model as a Proxy Model

You can designate a device model as a proxy model from the Information tab of the selected device model. Setting a device model up as a proxy model disables event creation for that model.

Note: When you have multiple proxy models in a Global Collection topology, you can collapse them to a single icon, merging all connections.


Follow these steps:

1. Select the device model to designate as a proxy model.
2. Click the Information tab, and expand the CA Spectrum Modeling Information subview.
3. Locate the 'Is a Proxy Model' setting, click 'set,' and select Yes.

Event creation is disabled for this model. The model now serves as a proxy model.


Cross-Landscape Fault Correlation Example

The following diagram provides an example of a network with multiple landscapes:




Landscape A, the local landscape, contains core routers, including Router R. Landscape B, the remote landscape, contains customer routers with the core Router R modeled a second time, this time as a proxy, as shown in the following diagram. This proxy model has its `IsEventCreationEnabled` attribute set to No. The device is being polled, but will not generate events or alarms.

Router R in Landscapes A and B:



If Router R goes down, as shown in the following diagram, Landscape B loses contact with the proxy and customer routers. However, only one red alarm is generated, in Landscape A. The proxy router stays green in Landscape B, where the alarms are suppressed because the proxy model's `IsEventCreationEnabled` attribute is set to No.

Alarm with Cross-Landscape Fault Correlation:



Configuring Port Status Monitoring

CA Spectrum provides the following methods of monitoring the status of ports:

Link Traps

Link traps allow you to monitor the status of ports without the cost of polling. However, traps are not always the most reliable notification mechanism of port status.

PollPortStatus

The PollPortStatus feature lets you poll the status of a port even if its connectivity is not modeled in CA Spectrum.

Live Pipes

Live Pipes let you turn on port status monitoring for individual links. This is a more reliable monitoring method than traps because CA Spectrum will periodically poll the status of the link (with an increased cost in performance). In addition, Live Pipes let you graphically verify which links are being monitored.

WA_Link Port Monitoring

WA_Link models automatically enable a live pipe for any connected ports.

NetworkLinkType

CA Spectrum automatically maintains the port-level attribute NetworkLinkType based on the model class of the two connected devices. The attribute lets you set up management policies based on the type of link a port is involved in.

Note: For more information about the NetworkLinkType attribute and management policies in general, see the *Policy Manager User Guide*.

The possible values for NetworkLinkType include the following:

- 0 = No Link
- 1 = Router Link
- 2 = Switch Link
- 3 = Shared Access Link
- 4 = End Station Link
- 5 = Wide Area Link
- 6 = Internal Link
- 7 = Unknown Link
- 8 = Network Cloud Link

If the connectivity of the port is not modeled, NetworkLinkType will be set to Unknown Link. When the connectivity of the port is modeled, the value of NetworkLinkType is maintained as described in the following table:

Connection	Attribute Value	Link Type
Router > Router	1	Router
Router > Switch Router	1	Router
Router > Switch	1	Router
Router > Hub	1	Router
Router > Workstation Server	1	Router
Switch Router > Switch Router	2	Switch
Switch Router > Switch	2	Switch
Switch Router > Hub	3	Shared Access
Switch Router > Workstation Server	4	End Station
Switch > Switch	2	Switch
Switch > Hub	3	Shared Access
Switch > Workstation Server	4	End Station

Connection	Attribute Value	Link Type
Hub > Hub	3	Shared Access
Hub > Workstation Server	4	End Station
Any port connected to a WA_Link model	5	Wide Area
Any backplane connecting inside a hardware device	6	Internal Link
EVPN Discovery is run and Provider_Clouds are created	8	Network Cloud Link

More information:

[Wide Area Link Monitoring](#) (see page 221)

[Live Pipes and Fault Management](#) (see page 224)

[Link Traps](#) (see page 218)

[PollPortStatus Feature](#) (see page 219)

Port Status Polling Criteria

In general, the status of a port is polled when the following criteria are met:

- The PollingStatus (0x1154f) of the port model must be TRUE.
- The Polling_Interval (0x10071) of the port model must be non-zero.
- The PollingStatus of the port's device model must be TRUE.
- Neither the port model nor the port's device model can be in maintenance mode; the isManaged (0x1295d) attribute for both must be set to TRUE.

If this criteria is met, port polling occurs with a frequency equal to the Polling_Interval setting.

However, either of the following two conditions override the default polling frequency:

- The port has been down since it was modeled in CA Spectrum and the Port Always Down Alarm Suppression attribute is set to Enabled.

Note: If the Port Always Down Alarm Suppression attribute is set to Disabled, the port will be polled as described above.

- The port is administratively down (that is, ifAdminStatus attribute is set to Down).

If these criteria are met, the polling frequency is reduced to once per hour (every 3600 seconds). Plus, all red alarms on the down ports are suppressed, and a gray condition is asserted. Administratively down ports remain brown.

More information:

[Port Status Monitoring Settings](#) (see page 226)

[Link Traps](#) (see page 218)

Port Status Events and Alarms

The port status monitoring engine uses the events and alarms listed in the following table to notify you of a change in status.

Event Description	Event ID	Alarm Description	Alarm ID	Port Condition Color
Port status good	0x10d10	N/A	N/A	Green
Port status bad	0x10d11	BAD LINK	0x1040a	Red
Port status disabled	0x10d12	LINK DISABLED	0x1040b	Brown
Port status unknown	0x10d13	LINK STATUS UNKNOWN	0x1040e	Gray
Port status unreachable	0x10d14	UNREACHABLE LINK	0x1040c	Gray
Port status initial	0x10d15	N/A	N/A	Blue
Port lower layer down	0x10d16	BAD LINK, BUT ALARM WAS SUPPRESSED	0x1040f	Gray
Port up, but linked with down port	0x10d17	LINK MAY NOT BE UP	0x10410	Gray
Port connected to down port or device	0x10d18	PORT ALARM SUPPRESSED	0x10411	Gray
Port status bad, but connected to WA_Link whose LinkFaultDisposition is LinkOnly	0x10d2d	PORT ALARM SUPPRESSED	0x10d2d	Gray

More information:

[Link Traps](#) (see page 218)

[Receiving Port Alarms](#) (see page 225)

Link Traps

Traps provide a means for network devices to let a management system know that a significant event has occurred on the network. Link Down and Link Up traps are perhaps the most important traps when it comes to port status monitoring. These traps tell the management system that a port has either become inoperable, or has come back up.

When CA Spectrum receives a Link Down trap, it polls the status of the corresponding port once to verify its status and generates one of the events and alarms listed in [Port Status Events and Alarms](#) (see page 217) on the affected port.

CA Spectrum generates a yellow alarm on the device model to allow easy access to vendor-specific trap data, however it no longer generates the trap-specific events and alarms on the affected port model.

Once a Link Down trap is received, CA Spectrum sets the OutstandingLinkDownTrap attribute on the port to TRUE. This will cause CA Spectrum to poll the status of the port regardless of the port status polling criteria. When CA Spectrum receives a Link Up trap for the port, or when the port's status is determined to be up based on a poll, the value of the OutstandingLinkDownTrap attribute is set to FALSE and polling will take place based on the value of the port status polling criteria. For more information about when a port is polled, see [Port Status Polling Criteria](#) (see page 216).

When all of the ports for which CA Spectrum has received a Link Down trap are back up, the yellow alarm on the device will be cleared.

You can use the following attributes to control how CA Spectrum handles link traps:

AlarmOnLinkDownIfTypes

This attribute contains a mapping of ifType values to a value which determines how to handle the trap for that particular ifType and model type (0 for never, 1 for always, and 2 for check admin). This can be customized in the MTE on a per-model type basis. When a port model is created, its AlarmOnLinkDownTrap (0x11fc2) attribute is automatically populated with the value which corresponds to its particular ifType.

ID: 0x1290f

AlarmOnLinkDownTrap

This attribute sets the alarm generation behavior for receipt of a LINK DOWN Trap Event. Possible settings are:

- Never (0) = Do not generate an alarm upon receipt of LINK DOWN trap
- Check Status (1) = Generate an alarm based on the current Admin Status (UP = Generate a Red alarm, otherwise generate a Brown alarm)

ID: 0x11fc2

AssertLinkDownAlarm

This attribute is used to determine if the yellow alarm should be generated on the device model. It is read from the port model for which CA Spectrum received the trap. This attribute is available from the port model's Attributes tab.

ID: 0x12957

More information:

[Configuring Port Status Monitoring](#) (see page 214)

Interface Trap Configuration

For many device models, you can configure the processing of link down traps received for individual port models through the port model's Interfaces tab, Component Detail view, Attributes tab. Here you can access the attributes that let you suppress link down alarms for the selected port model or its parent device model. Consult the CA Spectrum management module guide for the type of device you are interested in to see if that module supports such trap configuration.

You can also use Locator search to select a set of port models and then the Attribute Editor to update in bulk.

PollPortStatus Feature

The PollPortStatus feature lets you monitor the status of a port even if the port's connectivity is not modeled. The PollPortStatus attribute exists for both Device and Port models with a different attribute ID for each of the two model types. This lets you enable and disable port status polling at the device or port level. By default, PollPortStatus is set to TRUE at the device level and FALSE at the port level.

To reduce network traffic, SNMP reads for polled ports on the same device are grouped together into larger SNMP requests. This provides performance benefits that are most noticeable when many ports are polled by this method in a single SpectroSERVER.

Utilizing PollPortStatus to Watch a Connected Port's Status

The PollPortStatus attribute at the device level (0x12809) controls port status polling on a per-device basis. If TRUE, polling is enabled for that device. If FALSE, no ports will be polled, even if a port model's PollPortStatus attribute is TRUE. When changed to FALSE, alarms will be cleared for any port which is not involved in a live pipe.

The PollPortStatus attribute at the port level (0x1280a) controls polling for each port model. If this attribute is set to TRUE (and PollPortStatus for the device is also set to TRUE), the status of the port will be polled, and alarms will be generated if needed. When changed to FALSE, any alarm on the port will be cleared. The following table shows that a port's status will be polled only if PollPortStatus is TRUE for both a given port model and its device model.

Device Model's PollPortStatus Value	Port Model's PollPortStatus Value	Results
FALSE	FALSE	Port status is not polled for any port on device
TRUE	FALSE	Port status is not polled for this port on device
FALSE	TRUE	Port status is not polled for this port on device
TRUE	TRUE	Port status is polled for this port on device

CA Spectrum watches the polling interval to determine when to poll port status. When a port is polled, the port's status is determined and an appropriate alarm is generated (RED, BROWN, or GRAY) if necessary.

If a BAD LINK alarm is generated on a port (alarm code 0x1040a), and later polling on that port is disabled by changing the value of PollPortStatus to FALSE and disabling Live Pipes, an event will be generated to automatically clear the BAD LINK alarm.

The PollPortStatus attribute can be set to TRUE while the Live Pipes functionality is enabled. This will not cause redundant network traffic.

Enabling Port Status Polling

You can enable port status polling for all future models of a given type using the Model Type Editor (MTE), or on a current, per-model basis using the Command Line Interface (CLI). For example, use the MTE to set PollPortStatus to TRUE for both device and port model types. Then, when polled, interface models will generate appropriate alarms when needed. PollPortStatus can also be set at both the device and port level using the Global Attribute Editor.

Note: For information about using the CLI to enable or disable PollPortStatus for a single model, see the *Command Line Interface User Guide*.

Wide Area Link Monitoring

CA Spectrum polls any ports connected to a WA_Link model for status automatically. This polling is controlled by the PollingStatus and Polling_Interval of the WA_Link model.

If the PollingStatus of a WA_Link model is TRUE, and its Polling_Interval is non-zero, CA Spectrum automatically makes the pipes that are connected to a WA_Link "live" and sets the PollPortStatus for port models to TRUE. The live pipe lets you visually verify that CA Spectrum is monitoring the status of the connected ports.

If you disable the live pipe but leave the PollingStatus of the WA_Link set to TRUE, the Polling_Interval set to a non-zero number, and the ports' PollPortStatus set to TRUE, CA Spectrum continues to monitor the status of the ports.

Note: WA_Link models can only represent point-to-point connections, such as T1 and T3 lines. No more than two devices can be connected to it at a time.

LinkFaultDisposition

The LinkFaultDisposition setting provides control and flexibility over fault alarming. When a wide area connection goes down, alarms can be generated on ports and on the link model. You can set the LinkFaultDisposition (0x129e2) attribute from the WA_Link model Attributes tab.

LinkFaultDisposition can be set to one of the following three modes:

BothPortsAndLink

If set to BothPortsAndLink, alarms are generated on both the connected ports and on the link model. This is the default setting.

PortsOnly

If set to PortsOnly, only the connected ports are alarmed, and the WA_Link is suppressed.

LinkOnly


If set to LinkOnly, only the WA_Link model will be alarmed, and the ports will be suppressed.

More information:

[Suggested Port Fault Settings for Optimal Fault Notification](#) (see page 229)

Wide Area Link Monitoring Scenarios

Consider the sample WA_Link network topology shown in the following diagram:



The following tables illustrate two possible WA_Link monitoring scenarios for this topology.

Scenario 1: Link Goes Down, Device B Loses Contact

LinkFaultDisposition	Port 1 Condition	WA_Link Condition	Port 2 Condition
BothPortsAndLink	Red	Red	Gray
PortsOnly	Red	Gray	Gray
LinkOnly	Gray	Red	Gray

Note: If the Suppress Linked Port Alarms setting is set to TRUE, then the alarm on the upstream port will be suppressed, even if Link Fault Disposition is set to BothPortsAndLink. If Link Fault Disposition is set to Ports Only, and the Suppress Linked Port Alarms setting is set to TRUE, CA Spectrum generates an alarm on the WA_Link model.

Scenario 2: Link Goes Down, Device B Remains Reachable

LinkFaultDisposition	Port 1 Condition	WA_Link Condition	Port 2 Condition
BothPortsAndLink	Red	Red	Red
PortsOnly	Red	Gray	Red
LinkOnly	Gray	Red	Gray

Note: If the Suppress Linked Port Alarms setting is set to TRUE, then only one of the ports will be alarmed (red). The other port will be suppressed (gray). If Link Fault Disposition is set to Ports Only, and the Suppress Linked Port Alarms setting is set to TRUE, CA Spectrum generates an alarm on the WA_Link model.

Wide Area Link Modeling Best Practices

When you model a Wide Area Link, be sure to supply the IP address for the `Network_Address` parameter. As a best practice, supply values for the `WA_Link` model `Network_Address` and `Network_Mask` attributes based on the subnet of the connected router interfaces. In a proxy environment, the "real" and "proxy" links must have the same value for `Network_Address`.

CA Spectrum relies on the `Network_Address` (0x12d7f) attribute of a `WA_Link` to find duplicate `WA_Link` models and to collapse proxies into a single icon. The `Network_Address` is the only unique attribute that can be used to find duplicate `WA_Links`. However, this attribute is not automatically populated by CA Spectrum unless the `WA_Link` models are created by Discovery. When you manually create `WA_Links`, the `Network_Address` and `Network_Mask` attributes are not automatically populated, even when they are connected to router interface models with valid IP addresses.

In our testing, we have found that the Global Collection topology view does not collapse proxied `WA_Link` models correctly if the `Network_Address` parameter is not configured.

Set the `Network_Address` attribute of the `WA_Link` model to the network ID of the subnet to which the interface belongs. For example, a serial interface has an IP address of 10.253.9.2 and a subnet mask of 255.255.255.252. Set the `Network_Address` attribute of the `WA_Link` to 10.253.9.0 (10.253.9.2 with subnet mask 255.255.255.252).

In addition, do not simply draw pipes between the router and `WA_Link` icons. Proper `WA_Links` require a nested `WA_Segment`. In addition, an interface model on each router must be connected to the `WA_Segment` (and not to the `WA_Link`). This modeling paradigm enables CA Spectrum to establish fully resolved connections. It also enables the proper display of pipes in the Global Collection topology view.

Port Layer Alarm Suppression

Devices that support advanced network technologies, such as Frame Relay and Link Aggregation, have logical entries in the ifTable representing higher layer interfaces. CA Spectrum models these logical layers according to the ifStackTable. If you set the use_if_entity_stacking (0x12a83) attribute to TRUE on a device model in the Attributes tab, CA Spectrum attempts to use information from RFC2737 (Entity MIB) to model these logical layers if the ifStackTable method fails.

When a monitored higher layer port goes down (such as a Frame Relay DLCI, or logical trunk interface), CA Spectrum will query the statuses of all lower layer interfaces before alarming the port which went down. If all of the lower layer interfaces are down as well, CA Spectrum will suppress the higher layer interface alarm. A key example is a physical Frame Relay interface going down which has multiple circuits provisioned on it. All of the higher layer DLCI port models will be suppressed, and the single red alarm will exist on the physical Frame Relay interface.

Port Criticality

You can assign a relative importance value to port models using the port criticality (0x1290c) attribute. The criticality of a port is used in the Impact Severity calculations of an alarm on any port which may cause a network outage. You can also display the criticality of a port in the Alarms tab to allow prioritization of port alarms. The port criticality attribute can be set for an individual port from the port model's Attributes tab.

Live Pipes and Fault Management

Live Pipes functionality lets you turn on port status monitoring for individual links and display the status of a link by using status color indicators. A link is a connection between two devices that CA Spectrum has resolved to the port level. Live Pipes display a combined condition color from the two resolved ports representing each side of the link.

When a pipe is deleted, CA Spectrum removes all of its underlying associations (such as links_with and connects_to). If the pipe being deleted represents more than one link, CA Spectrum asks you to confirm the deletion.

More information:

[Configuring Port Status Monitoring](#) (see page 214)

Enable or Disable Live Pipes System-Wide

Live Pipes are enabled system-wide by default. Without Live Pipes enabled on a system-wide basis, enabling Live Pipes for individual links is not available.

To enable or disable Live Pipes system-wide

1. Expand the Live Pipes subview in the VNM model's Information tab.
2. Click 'set' in the Live Pipes field and select Enabled or Disabled from the drop-down list as desired.

More information:

[Live Pipes \(Links\)](#) (see page 24)

Enable or Disable Live Pipes on Individual Links

Live Pipes for individual links are disabled by default. All individual pipes are gold or silver until Live Pipes are enabled. When an individual live pipe is enabled, the `ok_to_poll` (0x11dd8) attribute is set to TRUE for both ports that are involved in the link. The status of the linked ports is monitored if `ok_to_poll` is TRUE and the port status polling conditions in the Port Status Polling Criteria section are met.

You can enable a live pipe for an individual link.

Follow these steps:

1. Right-click the link that you want to enable as a live pipe, and select 'Enable/Disable Live Links.'
2. Select the check box for the link to enable and click OK.

The Enable/Disable Live Links dialog closes. The link that you selected is enabled as a live pipe.

More information:

[Live Pipes \(Links\)](#) (see page 24)

Receiving Port Alarms

To receive alarms on a port model, the following conditions must be met:

- Live Pipes must be enabled system-wide
- Live Pipes for the link of interest must be enabled separately

If there is no model on the other side of the link, an alarm can still be generated when the status of the link changes by setting a watch on the MIB-II `ifOperStatus` attribute. When the `ifOperStatus` attribute returns a value other than 1, an alarm is generated. With this method, an alarm can still be generated even if the port is intentionally down (the `ifAdminStatus` attribute has been set to OFF).

You can set the default value of `ok_to_poll` in the MTE for any port model type. When a port connection is made, CA Spectrum will set both ports' `ok_to_poll` attributes to their MTE default values so that the pipe will automatically become live if the you want. When the connection is removed, the value for `ok_to_poll` remains at the MTE default value.

When CA Spectrum notices a change in the link's status, it will generate one of the events and alarms listed in [Port Status Events and Alarms](#) (see page 217) on the two ports involved in the link, and will also change the color of the live pipe to reflect its new status.

Port Status Monitoring Settings

The settings described in the following table let you to control the service of Live Pipes. These settings appear in the VNM model's Information tab, Live Pipes subview.

Note: The settings described in this section apply to port status monitoring throughout CA Spectrum, not just ports associated with Live Pipes.

Live Pipes

Setting this (global) option to Disabled turns off all pipes in the SpectroSERVER and no status polling will be performed for any ports associated with a pipe. However, any port which also has the `PollPortStatus` attribute set to TRUE will still be polled for status changes.

Attribute ID: 0x11df9

Alarm Linked Ports

Setting this option to TRUE will cause ports with a good status to have a gray alarm generated on it when linked with a bad or unreachable port.

Attribute ID: 0x11fbd

Suppress Linked Port Alarms

Setting this option to TRUE causes ports with a bad status to suppress their red alarms and generate a gray alarm if either the linked port or the connected device is bad or unreachable. Only one port in the link will be alarmed red. The other will be gray.

Attribute ID: 0x11fbe

Port Always Down Alarm Suppression

Enabling this option will cause CA Spectrum to suppress red alarms and assert a gray condition if a port has always been down since first being modeled in CA Spectrum.

Attribute ID: 0x12a03

Note: Any port status alarm will be automatically cleared by CA Spectrum if the port's `ok_to_poll` and `PollPortStatus` attributes are both set to FALSE.

More information:

[Port Status Polling Criteria](#) (see page 216)

[Suggested Port Fault Settings for Optimal Fault Notification](#) (see page 229)

[PollPortStatus Feature](#) (see page 219)

Monitoring Physical and Logical Connections

CA Spectrum can monitor the related physical connections of multilink bundles that are resolved to the logical connection via a Live Pipe. When Live Pipes are enabled on the logical connection, the physical ports on one side of the multilink bundle are polled. If Live Pipes are disabled, the polling on the associated physical ports are stopped.

Note: You may need to configure the `MultiLinkVirtualIfTypes` (0x12e3d) attribute to contain the list of `ifType` values for multilink bundles for which CA Spectrum should also poll the related physical connections. This attribute is pre-configured with the `pppMultilinkBundle` (108) `ifType`.

To avoid creating multiple alarms when a single physical connection goes down, CA Spectrum only polls the physical connections on one side of the multilink bundle. The side that is polled is determined by reading the `Criticality` (0x1290c) attribute for the logical connections and choosing the side with the higher value. If the values are equivalent, the connection with the lowest model handle is chosen. Also, the `NetworkLinkType` (0x12a79) attribute on the physical port models on the side that is polled is set to the same value of the `NetworkLinkType` attribute of the multilink virtual interface model.

The polling of the non-resolved physical ports on one side of the multilink bundle results in the following alarm behaviors:

- An alarm is generated on a physical port model whenever a physical connection that is part of a multilink bundle goes down.
- If the logical connection is down and at least one of the related physical connections is up, an alarm is generated on a multilink virtual interface model.
- If all of the physical connections are down and the logical connection also goes down, the logical interface is put into the suppressed (grey) state.


Note: The physical ports on the side that is not being polled remain in the green state at all times, unless some other type of monitoring has been enabled on them.

Examples: Monitor Physical and Logical Connections


In the following examples, two devices are connected through a multilink bundle of three pairs of physical port links. “O” represents a physical port and “o” represents a logical interface. Live Pipes are enabled on the logical connection between the multilink bundles. There are physical connections between O1a and O2a, O1b and O2b, and O1c and O2c, but these connections are not resolved in CA Spectrum.

When Live Pipes are enabled on the logical connection between the multilink bundles, the three physical ports on Device 1 (O1a, O1b, O1c) are also polled because o1’s Criticality (10) is greater than o2’s Criticality (5).

If the physical connection between O1a and O2a goes down, an alarm is generated on port O1a and port O2a remains green:



If all three physical connections are down, alarms are generated on ports O1a, O1b, and O1c, ports O2a, O2b, and O2c remain green, and logical connections o1 and o2 are put into a suppressed state:



Suggested Port Fault Settings for Optimal Fault Notification

In SPECTRUM 7.0 and later, the default settings for both Suppress Linked Port Alarms and Port Fault Correlation have changed. The default value of Suppress Linked Port Alarms has changed from FALSE to TRUE. This setting will suppress red alarms on ports that are connected to another down port or unreachable device. The default value of Port Fault Correlation has been changed from Management Neighbors Only to All Connected Ports. This eliminates the need for you to manually configure management neighbors before a fault occurs so that port fault correlation will work properly. If you have not previously used 'Management Neighbors' you will not need this setting.

Setting Suppress Linked Port Alarms to FALSE and Port Fault Correlation to Management Neighbors Only will approximate the fault notification behavior of SPECTRUM 6.6 with Service Pack 3.

CA recommends changing the default setting for Link Fault Disposition on WA_Link models. The default setting of BothPortsAndLink will result in multiple alarms if the link fails. Consider changing the setting to Link Only or Ports Only. Link Only is best in environments where the name of the WA_Link models or notes on these models is meaningful. Changing the setting to Ports Only may be better if fault notification consistency is most important. That is, regardless of the topology, you would prefer an alarm on a port model if there is a link failure.

More information:

[Port Status Monitoring Settings](#) (see page 226)

[Port Fault Correlation Options](#) (see page 204)

[LinkFaultDisposition](#) (see page 221)

[Example: Port Fault Correlation Scenario 2](#) (see page 207)

Device Criticality

The Device Criticality setting, accessible from a device model's Attributes tab, specifies the relative importance of the device within the network being modeled. When CA Spectrum loses contact with the device, this value is summed for the device itself and all of its downstream neighbors for which a gray condition is now being asserted (because their actual status cannot be determined). The aggregate device criticality value is displayed as the Impact Severity value of the associated alarm in the alarm's Impact tab. The default value for all devices is "1"; you can increase this value as desired depending on how important you consider the device to be; the higher the number, the more critical the device is to the network.

Configuring Fault Management for Pingables

Device fault isolation is faster and more reliable when device models have knowledge of each other as neighbors. When a fault occurs, each device model that is lost sends an ARE_YOU_DOWN action to all of its neighbor models. Depending on the answers that neighbor device models send, the lost models either turn gray or red.

Establish neighbor relations by creating a Connects_to association between a device model and the port model of another device. The act of pasting a device model on the interface of another device model adds each device model to the other's neighbor list.

Pingable models lack ports. As a result, in older versions of CA Spectrum, neighbor associations could not be established for Pingable models without the use of an inferred connector, like a Fanout. However, it is a best practice to create a relationship between models that is resolved directly. Placing both neighbor models in a Fanout means that fault resolution occurs indirectly. Fanouts and other inferred connector model types resolve faults differently and less directly during fault isolation.

Connect Pingable Models

You can connect Pingable models to each other in the topology view. Connect models to create a Connects_to association between them and receive more status information.

You can connect Pingable models using one of the following methods:

- Establish neighbors by drawing pipes between device models. Drawing a pipe between two Pingable models establishes a Connects_to association between the two models, making them neighbors.
- You can create a Connects_to association between two Pingable models using the CA Spectrum Command Line Interface. Use the following syntax:

```
./create association rel=Connects_to  
lmh=<model handle of Pingable A>  
rmh=<model handle of Pingable B>
```

When the Connects_to association is established, you see a gold pipe between the two connected models. When a fault occurs, each Pingable model sends the other an ARE_YOU_DOWN action.

Mapping Traps from Other Models to Pingable Models

You can map traps from multiple IP addresses to a single pingable model using the Command Line Interface (CLI) update command. You create the mapping by adding the IP addresses to the deviceIPAddressList (0x12a53) attribute on the pingable model.

Before you can specify the mapping, you must add the following option to the .vnmrc file if the option is not already included in the file:

```
enable_traps_for_pingables=TRUE
```

You can also remove mappings with CLI, and you can configure OneClick to display IP addresses mapped to Pingable models.

To map traps from other IP addresses to a pingable model

1. Connect to the SpectroSERVER with CLI.

Note: For more information about using CLI, see the *Command Line Interface User Guide*.

2. Invoke the update command:

```
./update
```

3. Add additional IP addresses to the deviceIPAddressList attribute (0x12a53) for the Pingable model you want to designate as a trap destination. The following example shows three IP addresses added to the attribute:

```
update mh=<pingable's mh> attr=0x12a53,iid=10.253.8.34,val=0
update mh=<pingable's mh> attr=0x12a53,iid=10.253.8.65,val=0
update mh=<pingable's mh> attr=0x12a53,iid=10.253.9.17,val=0
```

4. Verify that the IP addresses were added:

```
show attributes attr=0x12a53 mh=<pingable's mh>
```

Enable a Device IP Address View for a Pingable in OneClick

If the Device IP Address List category is not included under the Information tab for Pingables, complete the following procedure.

To enable a Device IP Address view for a Pingable in OneClick

1. Open the view-pingabledetails-config.xml file located in the following directory:

```
<$SPECROOT>/tomcat/webapps/spectrum/WEB-INF/topo/config/
```

2. Uncomment the following line:

```
<field-subview idref="devipaddrlist-subview-config"/>
```

3. Restart OneClick.

The Device IP Address List category appears under the Information tab for Pingable models. If you map IP addresses to a pingable model, the addresses appear in the list.

Remove an IP Address Mapping from a Pingable Model

When addressing schemes change in your environment, keep model information up to date. Use the CA Spectrum Command Line Interface (CLI) to remove an IP address mapping from a Pingable model.

Follow these steps:

1. Connect to the SpectroSERVER with CLI.

Note: See the *Command Line Interface User Guide* for more information about using CLI.

2. Invoke the update command:

```
./update
```

3. Remove the IP address from the deviceIPAddressList attribute (0x12a53).

The following example shows one IP address removed from the attribute:

```
update mh=pingable mh attr=0x12a53,iid=10.253.8.65,remove
```

4. Verify that the IP address was removed:

```
show attributes attr=0x12a53 mh=pingable mh
```

False Management Lost or Contact Lost Alarms

There is a SUN recommended security procedure which can increase the chance of a false Management Lost or Contact Lost alarm being generated. This security procedure modifies the ARP timeout on a Solaris host. This change in the ARP configuration will cause an increased delay in ARP messages being sent by the system, which will increase the overall time it takes for a SNMP packet to be sent and replied to. Since the SpectroSERVER starts the timeout timer as soon as it hands the SNMP request off to the operating system, the overall delay exceeds the settings in the SpectroSERVER.

One way to identify this behavior is to invoke the following command from a terminal window on the Solaris host:

```
arp -a | wc -l
```

This command counts the number of entries in the ARP table. The count will increase and then it will suddenly decrease back to the initial value. The time it takes for the decrease to occur depends on the current ARP timeout setting.

To avoid this, you must return the ARP timeout value to its default as described in this section.

To return the ARP timeout value to its default

1. Open the following file:

```
/etc/init.d/nddconfig
```

2. Look for an entry similar to the following:

```
ndd -set /dev/ip ip_ire_arp_interval 600000
```

This entry modifies the default ARP timeout value.

3. Remove this entry.
4. Depending on which version of Solaris you are operating in, type *one* of the following commands to update your system (you must run as root):

- Solaris (2.8):

```
ndd -set /dev/ip ip_ire_arp_interval 1200000
```

- Solaris (pre-2.8):

```
ndd -set /dev/ip ip_ire_flush_interval 1200000
```

5. Reboot your system to apply these changes.

The ARP timeout returns to the default value.

Chapter 8: Modeling and Managing SNMPv3 Devices

This section contains the following topics:

[SNMPv3 Support](#) (see page 235)

[Edit SNMP v3 Profiles Dialog](#) (see page 239)

[Manually Model an SNMPv3 Device](#) (see page 240)

[Troubleshoot SNMPv3 Communication Issues](#) (see page 251)

SNMPv3 Support

Note: SNMPv3 standards require a unique engineID for each SNMP entity (or engine). The SNMP engine/application must have its own unique engineID whether it is a manager or an agent. RFC 3414 and RFC 3418 are the official SNMPv3 standards. See the IETF website (<http://www.ietf.org/rfc.html>) for more information.

SNMPv3 support includes the following:

- Authentication
- Privacy
- 64-Bit Counters

CA Spectrum models and concurrently manages devices that support SNMPv1, SNMPv2c, and SNMPv3.

SNMPv3 Authentication

SNMPv3 provides the following levels of security: non-authenticated, authenticated, and authenticated with privacy. Authentication in SNMPv3 uses an encryption algorithm to determine if a message is from a valid source. CA Spectrum supports the SNMPv3 standard for the authentication of messages. You specify an authentication password for a device model when you create it.

When an SNMP packet is converted to SNMPv3, security parameters are added to the SNMPv3 packet that is sent to the device. The SNMPv3 agent on the device checks the authenticity of the message to verify that the packet came from an authorized source.

SNMPv3 data sent from the device to CA Spectrum also uses similar security parameters. CA Spectrum receives the packet and verifies its authenticity.

CA Spectrum supports the following encryption algorithms for authentication:

- MD5 (Message Digest Algorithm): Produces a 128-bit (16 byte) message digest. This encryption algorithm is the default. You can model a device configured to use MD5, using 'Authenticated' or 'Authenticated with Privacy.'
- SHA (Secure Hash Algorithm): Produces a 160-bit (20 byte) message digest.

CA Spectrum uses MD5 by default; however, you can specify a different authentication encryption algorithm by prepending it to the password in the SNMP community string.

More information:

[Specify a Privacy Encryption Algorithm on a Per-Model Basis](#) (see page 249)

Enable SNMPv3 Privacy

Privacy in SNMPv3 uses an encryption algorithm to encode the contents of an SNMPv3 packet to verify that it cannot be viewed by unauthorized entities when routed over the network. CA Spectrum supports the SNMPv3 standard for the encryption of messages. You specify a privacy password for a device model when you create it.

If configured properly, the SNMPv3 message is sent by CA Spectrum using the password to encrypt the message before it goes out onto the network. The destination device decrypts the data when it receives it. The return data sent from the device to CA Spectrum is also encrypted.

CA Spectrum supports the following encryption algorithms for privacy:

- DES: Data Encryption Standard (DES) is a 64-bit standard that encrypts and decrypts data.
- 3DES: Data Encryption Standard (DES) is a 64-bit standard that encrypts and decrypts data three times.
- AES: Advanced Encryption Standard (AES) is a 128-bit standard, cryptographic algorithm that encrypts and decrypts data.

CA Spectrum supports the use of DES by default. You can specify a different privacy encryption algorithm by prefixing it to the password in the SNMP community string.

Follow these steps:

1. In the Topology tab of the Contents panel, click  (Creates a new model by IP).
The Create Model by IP Address dialog opens.

2. Complete the fields as appropriate.

Network Address

Specifies the IPv4 or IPv6 address for the device you want to model.

DCM Timeout (ms)

Specifies the timeout between retry attempts (in milliseconds).

Default: 3000 milliseconds (3 seconds)

DCM Retry Count

Type the number of times that the DCM should attempt to send a request to a device that is not responding.

Agent Port

Specifies the SNMP agent port.

Default: 161

3. Select the SNMP v3 option in the SNMP Communications Options section.
The SNMP Community String field becomes disabled.
4. Click Profiles to create a new SNMPv3 security profile.
The Edit SNMP v3 Profiles dialog opens.
5. (Optional) To specify the 3DES privacy encryption algorithm, do the following:
 - a. Enter a name in the Profile Name field.
 - b. Enter the same data that has been configured for full MIB access on the device in the User ID field.
 - c. Select Authentication with Privacy from the Authentication Type drop-down list.
 - d. Enter the authentication password in the Authentication Password and Confirm Authentication Password field.
 - e. Enter the following in the Privacy Password and Confirm Privacy Password fields:

3DES^<privpassword>

- f. Click Add to update the Profiles list with the new profile you have created.
 - g. Click OK to save your changes and close the Edit SNMP v3 Profiles dialog.
6. (Optional) To specify the SHA authentication encryption algorithm and the 3DES privacy encryption algorithm, do the following:
 - a. Enter a name in the Profile Name field.
 - b. Enter the same data that has been configured for full MIB access on the device in the User ID field.
 - c. Select Authentication with Privacy from the Authentication Type drop-down list.
 - d. Enter the following in the Authentication Password and Confirm Authentication Password field:
`SHA^<authpassword>`
 - e. Enter the following in the Privacy Password and Confirm Privacy Password fields:
`3DES^<privpassword>`
 - f. Click Add to update the Profiles list with the new profile you have created.
 - g. Click OK to save your changes and close the Edit SNMP v3 Profiles dialog.
7. Select the Discover Connections check box, if appropriate.
8. Click OK in the Create Model By IP Address dialog to accept your selections and close the dialog.

The model of the SNMPv3 device appears in the Topology tab. The privacy encryption algorithm you specified appears in the SNMP Community String field of the CA Spectrum Modeling Information subview for the model.

Note: The Edit SNMP v3 Profiles dialog is also accessible by clicking Profiles in the Configuration tab in the Discovery Console.

More information:

[Specify a Privacy Encryption Algorithm on a Per-Model Basis](#) (see page 249)

64-Bit Counters

The SNMPv3 standard provides support for 64-bit counters. CA Spectrum can access 64-bit counter MIB variables for all SNMPv3 devices that comply with this standard.

SNMPv3 Support Issues

The following are some issues related to SNMPv3 support.

get-bulk Command

CA Spectrum support of SNMPv3 does not include the get-bulk command.

View Access Control Model (VACM)

CA Spectrum supports the VACM features of SNMPv3, however, VACM is not recommended. CA Spectrum has features that allow for secure access to devices. If you give CA Spectrum full view access to all device MIBs, you receive effective monitoring and management performance.

Performance and Capacity

High processing resources are required for CA Spectrum to effectively manage SNMPv3 devices. More overhead is consumed using the Authentication and Privacy features due to the time it takes to decrypt and authenticate each message.

This affects the number of device models that a SpectroSERVER can manage. Therefore, CA recommends that you only model devices using SNMPv3 that benefit from SNMPv3 support; model all other devices using SNMPv1.

SNMPv3 Security User Names on SpectroSERVER

You cannot use the same user name more than once for the three levels of SNMPv3 (non-authenticated, authenticated, and authenticated with privacy). For example, if you are using the user name “user1” for SNMPv3 level 1 non-authenticated, you cannot use that same user name for SNMPv3 level 2 authenticated or for SNMPv3 level 3 authenticated with privacy.

Modeling SNMPv3 Devices

Devices cannot be modeled using SNMPv3 if “:” or “/” appear in the SNMPv3 username, authorization password, or privacy password.

Edit SNMP v3 Profiles Dialog

The Edit SNMP v3 Profiles dialog can be accessed by clicking Profiles in the Configuration tab in the Discovery Console or from the Create Model dialogs.

The following is an illustration of the Edit SNMP v3 Profiles dialog:

Profile Name	User ID	Authentication Type
New v3 profile	group_operators	Authentication with Privacy
test_lab	QA_testers	Authentication with no Privacy
install_ops	site_engineers	No Authentication

Profile Name:

User Id:

Authentication Type:

Authentication Password:

Confirm Authentication Password:

Privacy Password:

Confirm Privacy Password:

Manually Model an SNMPv3 Device

You can manually model SNMPv3 devices and set up new profiles using the Create Model by IP functionality in CA Spectrum. You cannot model SNMPv3 devices using the Model by Model Type feature in CA Spectrum.

Note: When you discover SNMPv3 devices on Cisco switches with VLANs, you cannot use the community_string@VLAN_ID format to index bridging information for each VLAN. You must create contexts instead. For more information, see the *Cisco Device Management Guide*.

Follow these steps:

1. In the Topology tab of the Contents panel, click  (Creates a new model by IP).
The Create Model by IP Address dialog opens.
2. Complete the fields as appropriate.

Network Address

Specifies the IPv4 or IPv6 address for the device you want to model.

DCM Timeout (ms)

Specifies the timeout between retry attempts (in milliseconds).

Default: 3000 milliseconds (3 seconds)

DCM Retry Count

Specifies the number of times that the DCM should attempt to send a request to a device that is not responding.

Agent Port

Specifies the SNMP agent port.

Default: 161

3. Select the SNMP v3 option in the SNMP Communications Options section.
The SNMP Community String field becomes disabled.
4. Take *one* of the following steps:
 - Select an existing profile from the V3 Profile drop-down list and go to Step 6.
 - Click Profiles to create a new SNMPv3 security profile.
The Edit SNMP v3 Profiles dialog opens.
 - a. Enter a name in the Profile Name field.
 - b. Enter the same data that has been configured for full MIB access on the device in the User ID field.
 - c. Choose *one* of the following SNMPv3 standard security options from the Authentication Type drop-down list:
 - **No Authentication:** Data sent from the CA Spectrum host system to the SNMPv3 device is not encrypted or authenticated (go to Step 4e).
 - **Authentication with no Privacy:** Data sent from the CA Spectrum host system to the SNMPv3 device is authenticated but it is not encrypted. Enter the same data that has been configured for full MIB access on the device in the Authentication Password field. Confirm the password (go to Step 4e).

- **Authentication with Privacy:** Data sent from the CA Spectrum host system to the SNMPv3 device is both encrypted and authenticated. In the Authentication and Password fields, enter the same data that has been configured for full MIB access on the device (go to Step 4e).
 - a. Click Add to update the Profiles list with the new profile you have created.
 - b. Click OK to save your changes and close the Edit SNMP v3 Profiles dialog.
- 5. Select the Discover Connections check box, if appropriate.
- 6. Click OK in the Create Model By IP Address dialog to accept your selections and close the dialog.

The model of the SNMPv3 device appears in the Topology tab.

Important! If you modify the User ID field in the Edit SNMP v3 Profiles dialog after your model has connected, you will lose contact with the SNMPv3 device. To regain management of the device, right-click the device model in the Topology tab of the Contents panel, and click Reconfiguration, Reset SNMPv3 Authentication.

After you model SNMPv3 devices, you can model the rest of the network using Discovery. Discovery does not overwrite any of the devices that you have already modeled.

More information:

[Create Model Dialog](#) (see page 68)

Modeling an SNMPv3 Device Using a CA Spectrum Toolkit

You can use one of the CA Spectrum toolkits, for example, Modeling Gateway, to create a device model that supports SNMPv3. By default, MD5 is the authentication algorithm that is used and DES is the privacy algorithm that is used. The algorithms can be overridden by the '^' character. Use the following syntax when specifying the SNMP community string for the model.

For an SNMP community string that uses both privacy and authentication, use the following syntax:

```
#v3/P:authpassword:privpassword/userid
```

authpassword

Specifies the authentication password for the device.

privpassword

Specifies the privacy password for the device.

userid

Specifies the user ID for the device.

Example 1

```
#v3/P:myAuthPW:myPrivPW/myUserID
```

For an SNMP community string that uses a non-default privacy algorithm (3DES) and a default authentication algorithm, use the following syntax:

```
#v3/P:authpassword:3DES^privpassword/userid
```

authpassword

Specifies the authentication password for the device.

privpassword

Specifies the privacy password for the device.

userid

Specifies the user ID for the device.

Example 2

```
#v3/P:myAuthPW:3DES^myPrivPW/myUserID
```

For an SNMP community string that uses a non-default privacy algorithm (3DES) and a non-default authentication algorithm (SHA), use the following syntax:

```
#v3/P:SHA^authpassword:3DES^privpassword/userid
```

authpassword

Specifies the authentication password for the device.

privpassword

Specifies the privacy password for the device.

userid

Specifies the user ID for the device.

Example 3

```
#v3/P:SHA^myAuthPW:3DES^myPrivPW/myUserID
```

For an SNMP community string that uses authentication only, use the following syntax:

```
#v3/A:authpassword/userid
```

authpassword

Specifies the authentication password for the device.

userid

Specifies the user ID for the device.

Example 4

```
#v3/A:myAuthPW/myUserID
```

For an SNMP community string that uses a non-default authentication algorithm (SHA) and no privacy, use the following syntax:

```
#v3/A:SHA^authpassword/userid
```

authpassword

Specifies the authentication password for the device.

userid

Specifies the user ID for the device.

Example 5

```
#v3/A:SHA^myAuthPW/myUserID
```

For an SNMP community string that does not use authentication or privacy, use the following syntax:

```
#v3/N/userid
```

userid

Specifies the user ID for the device.

Example 6

```
#v3/N/myUserID
```

More information:

[Change Security Information for a Device Model](#) (see page 245)

Model an SNMP v2c Device Using a CA Spectrum Toolkit

To use one of the CA Spectrum toolkits to create a device model that supports SNMP v2c, use the following syntax when specifying the SNMP community string for the model:

```
#v2/<SNMP community string>
```

<SNMP community string>

Specifies the SNMP community string of the device.

Example:

```
#v2/mySNMPcommunitystring
```

Change Security Information for a Device Model

You can change security information for an existing SNMPv3 device model. You can also convert an SNMPv1 device model to an SNMPv3 device model. You must add the appropriate security information to the device model.

Follow these steps:

1. Select the model that you want to modify, and click the Information tab in the Component Detail panel.
2. Expand the CA Spectrum Modeling Information subview, and click set in the SNMP Community String field.
3. Modify the SNMP Community String using a syntax listed in [Modeling an SNMPv3 Device Using a CA Spectrum Toolkit](#) (see page 243) to create the appropriate string.

Note: For more information about using CLI commands, see the *Command Line Interface User Guide*.

Add Context Name Information

You can add the SNMPv3 context name value to be sent with SNMPv3 messages for a particular device.

Follow these steps:

1. Select the model that you want to modify, and click the Information tab in the Component Detail panel.
2. Expand the CA Spectrum Modeling Information subview, and click set in the SNMP Community String field.
3. Add the context name value to the SNMP Community String field. For example, if the current SNMP community string is:
`#v3/P:authPass:privPass/myuserid`
4. To insert a context name value of 'quark,' add 'quark' to the SNMP community string as follows:
`#v3/P:authPass:privPass/quark/myuserid`

Specify an Authentication Encryption Algorithm on a Per-Model Basis

CA Spectrum supports both MD5 and SHA authentication encryption, although MD5 is the default. You can specify the alternate encryption algorithm (SHA) by prepending it to the password in the SNMP community string. Prefixing the encryption algorithm on the SNMP community string for a particular device model overrides the default algorithm for that device model only.

To specify a privacy encryption algorithm on a per-model basis

1. In the Topology tab of the Contents panel, click  (Creates a new model by IP).
The Create Model by IP Address dialog opens.
2. Complete the fields as appropriate.

Network Address

Specifies the IPv4 or IPv6 address for the device you want to model.

DCM Timeout (ms)

Specifies the timeout between retry attempts (in milliseconds).

Default: 3000 milliseconds (3 seconds)

DCM Retry Count

Type the number of times that the DCM should attempt to send a request to a device that is not responding.

Agent Port

Specifies the SNMP agent port.

Default: 161

3. Select the SNMP v3 option in the SNMP Communications Options section.
The SNMP Community String field becomes disabled.
4. Click Profiles to create a new SNMPv3 security profile.
The Edit SNMP v3 Profiles dialog opens.
5. (Optional) To specify the SHA authentication encryption algorithm, do the following:
 - a. Enter a name in the Profile Name field.
 - b. Enter the same data that has been configured for full MIB access on the device in the User ID field.
 - c. Select Authentication with Privacy from the Authentication Type drop-down list.
 - d. Enter the following in the Authentication Password and Confirm Authentication Password field:
`SHA^<authpassword>`
 - e. Enter the privacy password in the Privacy Password and Confirm Privacy Password field.
 - f. Click Add to update the Profiles list with the new profile you have created.
 - g. Click OK to save your changes and close the Edit SNMP v3 Profiles dialog.

6. (Optional) To specify the SHA authentication encryption algorithm and the 3DES or AES privacy encryption algorithm, do the following:
 - a. Enter a name in the Profile Name field.
 - b. Enter the same data that has been configured for full MIB access on the device in the User ID field.
 - c. Select Authentication with Privacy from the Authentication Type drop-down list.
 - d. Enter the following in the Authentication Password and Confirm Authentication Password field:

`SHA^<authpassword>`
 - e. Enter the following in the Privacy Password and Confirm Privacy Password fields:

`[3DES|AES]^<privpassword>`
 - f. Click Add to update the Profiles list with the new profile you have created.
 - g. Click OK to save your changes and close the Edit SNMP v3 Profiles dialog.
7. Select the Discover Connections check box, if appropriate.
8. Click OK in the Create Model By IP Address dialog to accept your selections and close the dialog.

The model of the SNMPv3 device appears in the Topology tab. The authentication and privacy encryption algorithms you specified appear in the SNMP Community String field of the CA Spectrum Modeling Information subview for the model. You can also specify a privacy encryption algorithm or an authentication encryption algorithm by clicking set in the SNMP Community field of the CA Spectrum Modeling Information subview.

Note: The Edit SNMP v3 Profiles dialog is also accessible by clicking Profiles in the Configuration tab in the Discovery Console.

Change the Default Authentication Encryption Algorithm For All Device Models

To change the default authentication encryption algorithm for all device models, modify the .vnmrc file.

Follow these steps:

1. Navigate to the following directory:

`<$SPECROOT>/SS/`

2. Open the .vnmrc file with a text editor.
3. Locate the following line:

`snmpv3_auth_protocol=1`

Note: A value of "1" represents the default authentication encryption algorithm, MD5.

4. To modify the algorithm to use SHA as the default, change the parameter as follows:

`snmpv3_auth_protocol=2`

Specify a Privacy Encryption Algorithm on a Per-Model Basis

CA Spectrum supports DES, 3DES, and AES privacy encryption and uses DES by default. You can specify an alternate encryption algorithm by prefixing it to the password in the SNMP community string. Appending the encryption algorithm on the SNMP community string for a particular device model overrides the default algorithm for that device model only.

To specify a privacy encryption algorithm on a per-model basis

1. In the Topology tab of the Contents panel, click  (Creates a new model by IP). The Create Model by IP Address dialog opens.
2. Complete the fields as appropriate.

Network Address

Specifies the IPv4 or IPv6 address for the device you want to model.

DCM Timeout (ms)

Specifies the timeout between retry attempts (in milliseconds).

Default: 3000 milliseconds (3 seconds)

DCM Retry Count

Type the number of times that the DCM should attempt to send a request to a device that is not responding.

Agent Port

Specifies the SNMP agent port.

Default: 161

3. Select the SNMP v3 option in the SNMP Communications Options section.
The SNMP Community String field becomes disabled.
4. Click Profiles to create a new SNMPv3 security profile.
The Edit SNMP v3 Profiles dialog opens.
5. To specify the 3DES privacy encryption algorithm or the AES privacy encryption algorithm, do the following:
 - a. Enter a name in the Profile Name field.
 - b. Enter the same data that has been configured for full MIB access on the device in the User ID field.
 - c. Select Authentication with Privacy from the Authentication Type drop-down list.
 - d. Enter the authentication password in the Authentication Password and Confirm Authentication Password field.
 - e. Enter the following in the Privacy Password and Confirm Privacy Password fields:

`[3DES|AES]^<privpassword>`
 - f. Click Add to update the Profiles list with the new profile you have created.
 - g. Click OK to save your changes and close the Edit SNMP v3 Profiles dialog.
6. Select the Discover Connections check box, if appropriate.
7. Click OK in the Create Model By IP Address dialog to accept your selections and close the dialog.

The model of the SNMPv3 device appears in the Topology tab. The privacy encryption algorithm you specified appears in the SNMP Community String field of the CA Spectrum Modeling Information subview for the model. You can also specify a privacy encryption algorithm by clicking set in the SNMP Community field of the CA Spectrum Modeling Information subview.

Note: The Edit SNMP v3 Profiles dialog is also accessible by clicking Profiles in the Configuration tab in the Discovery Console.

Change the Default Privacy Encryption Algorithm For All Device Models

To change the default privacy encryption algorithm for all device models, you must modify the .vnmrc file.

To change the default privacy encryption algorithm for all device models

1. Go to the following directory:

```
<$SPECROOT>/SS/
```

2. Open the .vnmrc file with a text editor and locate the following line:

```
snmpv3_priv_protocol=1
```

Note: “1” represents the default privacy encryption algorithm, DES.

3. Depending on the privacy encryption algorithm you want to set as the default, modify the parameter as follows:

```
snmpv3_priv_protocol=2 (makes the default privacy AES)
```

```
snmpv3_priv_protocol=3 (makes the default privacy 3DES)
```

Troubleshoot SNMPv3 Communication Issues

An error message or alarm is displayed if CA Spectrum cannot communicate with an SNMPv3 device.

Consider the following:

Is the Device Model's Security Information Correct?

If you changed security information for a particular device model (see [Changing or Adding Security Information to a Device Model](#)), and the new information provided does not match the security information on the device, CA Spectrum generates an alarm indicating that it cannot contact the device using SNMP.

To troubleshoot this problem, update the security information for the device model to match the information on the device.

What Should I Do When CA Spectrum Loses SNMPv3 Contact with Cisco Routers After They Have Been Rebooted?

CA Spectrum can lose communication with Cisco devices, such as Cisco router models like 2621 v12.2 (IOS), 2517 v12.0 (IOS), or 2514 v12.2 (IOS).

SNMPv3 support includes a security feature named replay protection, which guards against SNMPv3 packet deciphering activities over the network. Replay protection checks the following two values on a device whenever a SNMP query is initiated:

- `snmpEngineBoots`: The number of times the device has rebooted.
- `snmpEngineTime`: The number of seconds since the `snmpEngineBoots` counter was last incremented.

CA Spectrum monitors these values for every device. When SNMP communication is properly occurring, CA Spectrum and a device are in sync with one another. If a device goes down, CA Spectrum receives the `snmpEngineTime` with the value of 0. CA Spectrum compares the `snmpEngineBoots` value and, if it has incremented, communication resumes. If the `snmpEngineBoots` value has not incremented, then CA Spectrum does not resume communication.

This problem is due to a Cisco IOS firmware bug that will not increment the boot count causing SDManager to stop communication.

To avoid this performance problem, upgrade these routers with the latest Cisco IOS firmware. See <http://www.cisco.com> for details.

Note: For more information about replay protection, see RFC 3414, section 2.2, Replay Protection.

Chapter 9: CA Spectrum Intelligence

This section contains the following topics:

[Inductive Modeling Technology](#) (see page 253)

[Static Configuration of Device Models](#) (see page 253)

[Dynamic Configuration of Device Models](#) (see page 254)

[Condition Versus Rollup Condition](#) (see page 256)

[Fault Isolation](#) (see page 264)

[Duplicate Addresses](#) (see page 272)

[Manually Clear Duplicate Addresses](#) (see page 274)

[Automatic Naming and Addressing](#) (see page 274)

[Detection of Firmware Problems](#) (see page 275)

[Interface Intelligence](#) (see page 275)

Inductive Modeling Technology

CA Spectrum comes with Inductive Modeling Technology™ (IMT), a patented technology that consists of a suite of intelligence circuits which work with the VNM to help configure, manage, and monitor your network.

Static Configuration of Device Models

Many network devices are configured during their manufacturing process and are difficult to modify later. For CA Spectrum modeling purposes, these devices are considered to have a static configuration--once CA Spectrum intelligence models these devices, they are not reconfigured. CA Spectrum creates models for the ports, matching the types of port models to the types of ports on the device, such as T1 or Ethernet. For each port model that is created, an association is established between the port and the device using the HASPART relation and may be viewed within the device model's Interfaces tab. When the model is destroyed, all of the port models associated with the device are also destroyed.

Note: For more information about the Interfaces tab, see the *Operator Guide*.

Dynamic Configuration of Device Models

Some network devices can be configured dynamically by removing boards and installing new boards without removing the device from the network. CA Spectrum intelligence provides for automatic modeling of these devices and their connections upon their creation and then performs verification and remodeling if necessary after each VNM polling cycle. Thus, CA Spectrum continuously monitors and changes these models to match the actual device on the network.

Whenever a model is created, SpectroSERVER polls the device and creates an appropriate configuration, including the number, type, and order of modules and ports. For each device model created, a relation exists between that model and the parent device via the HASPART model type relation rule. This relation also exists between the boards and any ports on the boards. After each polling cycle, SpectroSERVER re-examines and, if necessary, changes the parent model and its related models to match changes to the device configuration.

Whenever you add a new board to a dynamically configured device, CA Spectrum creates a model to represent that board and each of the ports on the board. If a board model is destroyed, all of the port models that form part of the board are also destroyed.

Pulled Board List

Creating and destroying models is time consuming. When you remove a board from the device, CA Spectrum does not destroy the board model, but rather keeps a copy of the board model in a “pulled board list.” The board model’s HASPART relation to the hub model is removed. CA Spectrum reassociates this model to the device if you reinstall the old board. If you add a new board to replace the old board, CA Spectrum associates the new board model to the device and the old board model is placed in the Lost and Found view. If you remove a board model from the pulled board list, the board model is removed from the Lost and Found view and no longer exists.

The following are general pulled board list attributes:

Max_Pulled_Bd_Cnt

Specifies the maximum number of models allowed to exist in the pulled board list. When this value is exceeded, the oldest model is removed from the list.

Pulled_Bd_Cnt

Specifies the current number of models in the pulled board list.

Pulled_Bd_List

Contains a list of board models that have been pulled from a dynamically configured device. When a board is reinstalled in the device, CA Spectrum removes the board model from the Pulled_Bd_List.

Note: The Pulled_Bd_List is not readable by the user.

Router Reconfiguration Events

Router reconfiguration actually involves two separate processes: interface reconfiguration, which helps ensure the device's interfaces are properly modeled, and device discovery, which helps ensure proper modeling of other devices, LANs, and so on that are connected to those interfaces. Depending on whether both or either one of these processes occurs, CA Spectrum generates one of the following events to help you keep track of the configuration changes:

ROUTER_RECONFIG_EVENT (0x1001c)

This event is generated when a device is reconfigured and both interface reconfiguration and device discovery occur.

INTERFACE_RECONFIG_EVENT (0x1001d)

This event is generated for a device whenever interface reconfiguration occurs.

DEVICE_DISCOVERY_EVENT (0x1001e)

This event is generated for a device whenever device discovery occurs such as when connections off the device's interfaces are being rediscovered.

Condition Versus Rollup Condition

CA Spectrum provides intelligence circuits that let you see changes in your network devices and their performance by simply glancing at the icons that represent them. The icons use color to indicate two different types of status: Condition and Rollup Condition. Condition reflects the contact and alarm status of the modeled device represented by the icon. Rollup Condition is the *composite* status of models that are “children” of the model represented by the icon. (Child models are related to parent models through the “collects” relation in the Topology hierarchy and through the “contains” relation in the Location hierarchy.) The Rollup Condition generally changes as you move up in the hierarchy, because at each level it reflects the blending of the Rollup Conditions from a greater number of individual models.

The location of the Condition and Rollup Condition colors varies according to the type of icon. For device and topology (LAN) icons, Condition is displayed in the diagnostic double-click zone and the Rollup Condition is displayed in the down-arrow double-click zone for the icon. The circle at the base of location model icons displays either the Condition or the Rollup Condition, whichever is more critical.

Attributes Determining Condition and Rollup Condition

There is a unique set of attributes that are related to Condition and Rollup Condition. Their values are used in determining Condition and Rollup Condition for models:

Condition

The Condition attribute value reflects the contact status as well as any more specific alarm in effect for a device model. This value determines the Condition color on topology and location icons as explained in the following table:

Contact Status	Condition	Color	Description
Initial	Initial	Blue	Either the model has not yet established contact with the device it represents, or it represents an insignificant device with which contact has been lost.
Established	Normal	Green	The model has successfully established contact with the device it represents, and the device is functioning normally.
Established	Minor	Yellow	This is the first level of marginal operation. Either the model has successfully established contact with the device it represents but there is an abnormal condition that does not affect overall network operation (perhaps a module has been removed from the device), or the IP address assigned to this model was already assigned to another model.

Contact Status	Condition	Color	Description
Lost	Major	Orange	This is the second level of marginal operation. The management agent on the device has failed and is not responding to any communication from CA Spectrum but the device is still relaying data to its downstream neighbors. This condition occurs only on data-relay type devices such as hubs and is typical of a firmware failure.
Lost	Critical	Red	This condition indicates a total failure of the device and requires management's attention to repair or replace it.
Lost	Suppressed (Unknown)	Gray	Contact has been lost with this device <i>and</i> with a device that is upstream from this device (for example, between this device and CA Spectrum), thus the actual condition of this device is unknown and alarms for the model representing it are suppressed. The gray condition color is also displayed for all models that are downstream from this device. All adjacent (directly connected) models, whether upstream or downstream, will have a contact status of "Lost."
Lost	Maintenance	Brown	CA Spectrum cannot contact the device because the model has been placed into maintenance mode.

Condition_Value

Specifies a numeric value that represents a model's overall condition. This value is passed to a parent model and included in the composite condition. The model's overall condition is either the condition or the rollup condition, whichever is more severe.

Note: The condition value indirectly receives the value of the administrator-defined significance level corresponding to a model's overall condition.

Composite_Condition

The sum of all condition values for models that are contained by a location model or collected by a topology model.

Rollup Condition

CA Spectrum computes the rollup condition attribute value using the administrator-defined rollup threshold and composite condition. The resulting attribute value determines the color that is displayed to indicate the overall condition of models that are contained by a location model or collected by a topology model. The possible colors are:

Green

The value of the composite condition attribute for this model's children is less than the yellow (rollup condition) threshold.

Yellow

The composite condition value for this model's children equals or exceeds the yellow threshold but is less than the orange threshold.

Orange

The composite condition value for this model's children equals or exceeds the orange threshold but is less than the red threshold.

Red

The composite condition value for this model's children equals or exceeds the red threshold.

Condition and Rollup Condition Sensitivity

The following two attributes serve as parameters that can be used to emphasize or diminish the impact on rollup condition from the condition values of particular models. By adjusting these attribute values you can control when the rollup condition color changes.

Rollup Thresholds

The rollup thresholds are the three attributes that control the rollup condition color (yellow, orange, and red) for a model. Rollup thresholds are administrator-defined values that are entered on a model-by-model basis. The composite condition value received from the model's children is compared with these attributes to determine a rollup condition color. For example, if a model's composite condition value is equal to or greater than its orange threshold (but less than its red threshold), the model's rollup condition color is orange.

The default values for rollup thresholds are:

- Yellow Threshold = 3
- Orange Threshold = 6
- Red Threshold = 10

Significance Level

The significance level attributes define the numeric value for yellow, orange, and red conditions and rollup conditions. Like rollup thresholds, significance level values are administrator-defined values, and are entered on a model-by-model basis.


Significance level field labels begin with the words “value when.” The default significance level values are:

- Value_When_Yellow = 1
- Value_When_Orange = 3
- Value_When_Red = 7

Typically, models (devices) are divided into two classes, “significant” or “insignificant.” A significant device is any device that requires an administrator’s attention for proper network operation. Insignificant devices are typically end-point devices, such as a PC or workstation. Insignificant devices usually toggle between green and blue (Condition Value = 0). Significant models can be made insignificant by changing their Value_When_Red attribute value to 0 (zero).


Rollup Condition Flow

An overview of the rollup condition process is illustrated in the following diagram. Read the flow from bottom to top, but keep in mind that it shows a single path in the propagation of rollup condition and that there may be many children passing condition values to a parent model.



Example of Rollup Condition Propagation

The following diagram illustrates the propagation of a rollup condition in the Topology hierarchy.



This example depicts two layers of a Topology hierarchy. The example assumes the use of default Rollup Thresholds and Significance Levels. At the lowest level in the figure there are five devices: two hubs, a router, and two end-point devices. These are contained by TWLAN, a LAN of type 802_3_LAN (as are the two other LANs: FINANCE and ENGRG). The network group model named OVERALL_NET collects these three LAN models.

The following Rollup Conditions and Conditions, at lower levels, determine the top-level Rollup Condition for the network group model OVERALL_NET:

Devices Collected by TWLAN

Hub#1

Condition = Green
Rollup Condition = Red
Condition Value = 7

Hub#2

Condition = Green
Rollup Condition = Orange
Condition Value = 3

Router#1

Condition = Green
Rollup Condition = Yellow
Condition Value = 1

End-Point Devices PC#1 & PC#2

Condition = Green
Rollup Condition = Green
Condition Value = 0

LAN Named FINANCE

Condition = Green
Rollup Condition = Green
Condition Value = 0

LAN Named ENGRG

Condition = Green
Rollup Condition = Yellow
Condition Value = 1

Example Rollup Condition Process

Every model within a Topology view receives its Collects relation from the model that it is collected by. Therefore, all models contribute to the rollup condition of the network group model. The following steps provide a detailed flow of condition values that contribute to the rollup condition for the model named OVERALL_NET shown in the previous example diagram.

1. Determine the Composite Condition for the TWLAN network group model. Composite Condition is the sum of the collected models' Condition Values. In this case, the device models have Condition Values of:

“Orange” Condition hub model has a Condition Value of 3.

“Red” Condition hub model has a Condition Value of 7.

“Green” Condition PC#1 model has a Condition Value of 0.

“Yellow” Condition router model has a Condition Value of 1.

“Green” Condition PC#2 model has a Condition Value of 0.

Therefore, for the TWLAN model:

Composite Condition = $(3 + 7 + 0 + 1 + 0) = 11$

2. Determine the Rollup Condition for TWLAN. In this case:

Composite Value = 11

TWLAN Yellow Threshold = 3

TWLAN Orange Threshold = 6

TWLAN Red Threshold = 10

Composite Value > Red Threshold

Therefore:

Rollup Condition for TWLAN = Red

3. Assign Significance Levels to TWLAN Condition and Rollup Condition. In this case, Significance Levels are:

Value When Yellow = 1

Value When Orange = 3

Value When Red = 7

Therefore:

Rollup Condition = Red Condition = 7

4. Set Condition Value for TWLAN model. In this case:

Rollup Condition more severe than Condition

Therefore:

Condition Value = Rollup Condition Significance Level = 7

The three network models TWLAN, ENGRG, and FINANCE pass their Condition Values up to the network group model OVERALL_NET. Changes in the Condition or Rollup Condition for the device models at lower levels can produce changes in the topology models further up in the Topology hierarchy. The Rollup Conditions for these three networks produce the following Rollup Condition for OVERALL_NET.

5. Determine the Composite Condition for the OVERALL_NET network group model. Composite Condition is the sum of the collected models' Condition Values. In this case, the network models have Condition Values of:

Red Condition TWLAN model has a Condition Value of 7.

Green Condition FINANCE model has a Condition Value of 0.

Yellow Condition ENGRG model has a Condition Value of 1.

Therefore, for the TWLAN model:

Composite Condition = $(7 + 0 + 1) = 8$

6. Determine the Rollup Condition for OVERALL_NET. In this case:

Composite Value = 8

Yellow Threshold = 3

Orange Threshold = 6

Red Threshold = 10

Composite Value > Orange Threshold, but < Red Threshold

Therefore:

Rollup Condition for OVERALL_NET = Orange

Fault Isolation

Fault Management is one of the key requirements of network management. A fault is different from an error because it is an abnormal condition that requires management attention and repair. Problems that give rise to a fault could be caused by bad firmware, bad hardware, or a bad network. Each of these problems requires a different response from the network manager. Thus the goal is to determine the exact location of the fault and to get the attention of the network administrators as quickly as possible.

CA Spectrum intelligence has the capability of isolating a network problem to the most probable faulty component. To speed up fault isolation and to reduce unnecessary traffic, two actions occur:

Are-You-Down Action

Upon losing contact with the device it represents, a model sends the Are-You-Down action to all of its neighbors to determine its own condition. If all of the neighbors return a response of TRUE, the model's condition color will turn gray (meaning "my device might be down, but it is impossible to tell because all the neighbors are down"). However, if any of the neighbors return a response of FALSE, the model's condition color will turn red (meaning "my device must be down, because one of the neighbors is up").

Are-You-Up Action

Upon re-establishing contact with the device it represents, a model sends the Are-You-Up action to its neighbors to speed up the fault isolation. Upon receiving this action, each neighbor will return TRUE if it has an established contact status. If the model's contact status is lost, and the next-time-to-poll is more than 60 seconds, then the model pings the device for quicker fault isolation.

Every time a model's status changes, or the information available to CA Spectrum changes, a new assessment occurs. CA Spectrum intelligence keeps the topology presentation as current and as accurate as possible, but it depends on correct modeling to accurately assess contact status and determine device failures on the network. Correct modeling includes placing your VNM model in proper relation to the other models that represent your network; it must have a resolved connection in the Topology view of a model that represents a device to which the VNM host is actually connected. When the VNM model is properly connected and CA Spectrum loses contact with a model, the icon representing that model displays a condition color of Gray, Orange, or Red, which helps the network administrators to locate the faults immediately.

How Model Category Affects Contact Status

Each fault is associated with a particular condition, which is represented by a particular color that displays on the icon representing the model where the fault occurs. The condition color reflects both the contact status and the alarm status of the model. However, the contact status and condition color asserted for a model also depend upon which of the following categories a model belongs to. The following list summarizes how the categories to which a model and its neighbors belong influences its contact status and condition color.

Significant Device Models

Any device that requires an administrator's attention for the smooth operation of the network is called a significant device. To change an insignificant model into a significant model change the value of the attribute Value_When_Red (0x1000e) to 7.

Insignificant Device Models

An insignificant device such as an end user PC toggles between Blue and Green contact states and does not generate alarms or event messages to get the attention of the administrator. To change a significant model into an insignificant model change the value of the attribute Value_When_Red (0x1000e) to 0.

Inferred Connectors

These are dumb models that do not poll, but that keep track of a list of their Data Relay neighbors. Possible inferred connectors are: WA_Segment, Fanout, and so on. CA Spectrum automatically enables Live Pipes for all ports connected to a WA_Segment.

Note: CA Spectrum intelligence does not expect Fanout models to be connected to each other; thus this configuration results in inaccurate contact status displays. If two Fanouts are connected to each other and each of them is in turn connected to a device with a green contact status, the Fanouts nonetheless turn red. If two Fanouts are connected to each other with no other devices connected to either one, both Fanouts turn gray.

Shared Media Link

The Shared Media Link is a specialized inferred connector. These models are similar to Fanouts, but the fault management works differently. Unlike a Fanout model, the Shared Media Link model condition is based on configured threshold values.

Example: If the critical threshold is set to 80, the Shared Media Link turns red when it loses contact with 80 percent of the downstream models.


Composite and Discrete Topology Models

The contact status of LAN, LAN 802.3, LAN 802.5 and so on models is determined by the contact status of its collected children. A LAN model with lost contact status will turn either red or gray, depending on the condition of its collected models.

Wide Area Links

Wide Area Links (WA_Links) are modeled in conjunction with wide area segment (WA_Segment) models. This allows for proper rollup of the Wide Area Link condition. WA_Link models can only represent point-to-point connections, such as T1 and T3 lines, and there can be no more than two devices connected to it at a time. Also, you must connect the WA_Segment model to the correct port of the device models.

Note: WA_Link models can accommodate only one WA_Segment model. If you attempt to paste more than one WA_Segment model into a WA_Link model's Topology view, the second one will be destroyed immediately and an alarm will be generated.



Wide Area Segments

WA_Segments poll the InternalPortLinkStatus (IPLS) attribute of each interface model which Connects_To the WA_Segment. This is an active poll, meaning that the IPLS of each connected interface is read at every polling interval rather than simply watched for a change in the attribute. Therefore, CA Spectrum does not have to lose contact with one of the connected routers for a fault isolation alarm to be generated on a WA_Link.

The polling of the connected ports' IPLS will be regulated by the WA_Link model's Polling_Interval and PollingStatus attributes. When the Polling_Interval changes to zero (0) or PollingStatus goes to FALSE, polling of the connected port's IPLS is stopped.

If one of the connected interfaces has an IPLS of BAD (for example, Admin Status is ON, but Oper Status is OFF), then the WA_Segment's Contact_Status is set to 'lost' and the WA_Segment turns gray. The WA_Link turns red.

If one of the connected interfaces has an IPLS of 'disabled' (for example, Admin Status is OFF), then the WA_Segment's Contact_Status is set to 'lost' and the WA_Segment turns gray. The WA_Link turns orange. This is because the alarm must be severe enough to be viewed in the Alarms tab, but it is not a "Contact Lost" alarm.

If the DISABLED interface causes CA Spectrum to lose contact with the remote router, then the WA_Link turns red. This is the regular InferConnector-type fault isolation working.

Model Category	Connected Models (Neighbors)	Condition Color
Significant Devices (Modeling Hub-types only)	connected to a VNM...	turn Red after losing contact
Significant Devices	with no connections to other models (a zero connector count)...	
Significant Devices	connected to an established Data Relay neighbor...	
Composite and Discrete Topologies	in which all of the collected children have a lost contact status and at least one of those collected children is Red...	
Inferred Connectors	where the fanout model has lost contact but one of its neighbors is good and the associated port has bad port link status, then it...	
Significant Devices, Inferred Connectors, and WA_Links	where all neighbors have also lost contact status...	turn Gray after losing contact.

Model Category	Connected Models (Neighbors)	Condition Color
Composite and Discrete Topologies	in which all ocs and none of those collected children are Red...	
Significant Devices (Modeling Hub-types only)	connected to an end-point neighbor (such as a PC) that has established contact status...	turn Orange after losing contact.
WA_Links	WA_Segment (or fanout) is good and one of the routers is lost then...	
Significant Devices	connected to a model with an Established contact status...	turn Green.
Composite/Discrete Topologies and WA_Links	in which any of the collected children has established contact status, then the LAN will also...	
Inferred Connectors	connected to a model with an established contact status where at least one of its neighbors is <i>Good</i> and its associated port (port connected to the Fanout) status is <i>Good</i> ...	
Significant and Insignificant Devices	not yet connected to other devices...	turn Blue
Composite/Discrete Topologies and WA_Links	when all collected children of a LAN have initial contact status, then the LAN will also have the initial contact status...	


Fault Isolation Examples

The following examples illustrate how CA Spectrum fault isolation operates with various network configurations and problem scenarios.

Example: Proactive Fault Isolation

This example demonstrates that fault isolation is a proactive mechanism which does not depend upon polling all of the connected models.

Consider a simple network topology as shown in the following diagram. The device H1 is connected to the VNM model. Devices H1, H2, and H3 poll every 3 minutes. H4 polls every 5 minutes. The PC polls every 30 minutes.



Assume H2 is BAD. As a result H2 turns red, H4 turns gray, PC (insignificant model) turns blue, while H1 and H3 remain green.


Fault isolation is initiated as soon as H2, H4 or PC polls. If H4 is lost, it sends an Are-You-Down action to H2. If H2 is lost by then, it sends TRUE to H4, otherwise it pings itself and then sends the response to H4. This causes H4 to turn gray.

Now H2 is lost, and it sends Are-You-Down action to H1. Since H1 is established, H2 has to decide between orange and red conditions. H2 pings PC. Since PC cannot respond H2 will turn red. The ping from H2 puts PC in a lost state. Since PC is an insignificant device it will turn blue.

Example: Modeling a Fanout

This example demonstrates fault isolation when modeling a fanout.

Assume the fanout is red and D2, D3, and PC are gray. The following diagram illustrates this scenario.



The fanout registers a watch on D1's contact status. If D1 goes down, the fanout turns Gray as a result of the watch trigger.

When D3 eventually polls successfully, D3 will have an established contact status and turn Green. D3 then sends an Are-You-Up action to the fanout. The fanout reads device P3's (D3's port connection to the fanout) internal link port status. Assuming the port has a good status, the watch is cleared and the fanout turns Green with an established contact status. This means that as long as P1 (D1's port connection to the fanout) has good internal link port status, the contact status of the inferred connector will remain good.


What if D2 goes bad? D2 will lose its contact status and sends an Are-You-Down action to the fanout. The fanout will ping D1, and finds D1 to be good. The intelligence then examines the status of P1. Assuming Link-Status of P1 is good, the fanout will return FALSE to model D2. This causes D2 to turn Red.

What if P1 is bad? This is the same case as disconnecting the network connection to the fanout. If D3 polls first, it will lose its contact status and send an Are-You-Down action to the fanout. The fanout will ping D1 as finds it as a good neighbor. fanout then reads the internal-port-link-status of the port P1. Since P1 is bad, the fanout will lose its contact status and turns Red. The fanout will return TRUE to the model D3. This causes D3 to turn Gray. D2 will also turn Gray in the same way as D3. PC being the insignificant device will turn Blue immediately after losing its contact status.

Example: Redundant Paths Fault Isolation

This example shows how CA Spectrum manages devices using redundant paths if a link is shut down administratively (i.e., admin-status equals *down*).

The following diagram depicts a network with redundant WA Links. Here VNM manages Rtr3 through link WL-1 and Rtr2 using link WL-2. Assume that the network administrator shuts down the WL-1 link. This causes WL-1 to turn gray. Rtr3 will turn red because VNM cannot talk to it through WL-1. The redundancy intelligence of Rtr3 will modify its agent address, so that VNM can talk to it using links WL-2 and WL-3. This causes Rtr3 to turn green again. The link WL-1 will still have the gray condition.




Example: Inferred Connector Fault Isolation

This example demonstrates that fault isolation for an Inferred Connector requires specific modeling. Assume that two routing devices, Rtr1 and Rtr2, are connected at both ends of the WA_Link and that their ports are P1 and P2 respectively.

WA_Link models need to be associated with a WA_Segment (or fanout) model through the Collects relation to enable the proper rollup of the WA_Link condition. The devices at either end of the WA_Link need to be connected to the WA_Segment collected by the WA_Link model. You do this by navigating into the device's Device Topology view and resolving the WA_Segment off-page reference icon to the appropriate port. You can view the connections by navigating into the WA_Segment's view.

This cross-connection is very important for fault isolation to work, as shown in the following diagram.



Assume P1 is the port on Rtr1 and P2 is the port on Rtr2. The routers connected to the WA_Segment will cause it to behave as described in the following table. Note that the port link status becomes important in determining the status of the WA_Link only when both routers are “contact established.”

Rtr1	Rtr2	WA_Link
Initial	Initial	Blue
Established	Lost	Red
Lost	Lost	Gray
Established	Established	Check Port States*

* If both Rtr1 and Rtr2 have a contact status of *established* then the port status of P1 and P2 will determine the condition of the WA_Link. If any port is BAD, the WA_Link will be RED. If any port is DISABLED, the WA_Link will be ORANGE. Otherwise, the WA_Link will be GREEN.

Duplicate Addresses

CA Spectrum intelligence automatically detects when duplicate IP addresses are entered in the SpectroSERVER database. Although some devices are allowed to have a duplicate IP address, Cabletron hub devices should be configured with only one IP address per device.

CA Spectrum can model different devices that share IPs, provided that each device has at least one IP that is unique to that device. This modeling policy accommodates certain networking technologies such as load balancing that create identical IP addresses across a range of devices. Devices that share some interface IP addresses can be modeled manually or by using Discovery. However, devices that share all their interface IP addresses in common cannot be modeled manually or by using Discovery.

Alarm condition colors warn you of duplication, as shown here:

Same MAC Address & Different IP Address

This alarm occurs when there are two or more models with the same MAC address and at least one model with a different IP address.

Color: Yellow

Same IP Address & Different MAC Address

This alarm occurs when there are two or more models with the same IP address and at least one model with a different MAC address.

Color: Orange

Same IP Address & Same MAC Address

This alarm occurs when there are two or more models with the same IP address and the same MAC address (duplicate addresses).

Color: Yellow

Duplicate MAC Address

The special case alarm for duplicate models where at least one of the models does not have an IP address. Only a Physical_Address model type can have this characteristic.

Note: Even if the MAC address for two device models is identical, this alarm occurs only when the MAC address of every interface of the two devices is the same.

Color: Yellow

To get these alarms, a model type needs both the MAC address and the IP address. For example, the model types Pingable and PhysicalAddress do not have both addresses, so you will not see these alarms.

Manually Clear Duplicate Addresses

You can clear duplicate address alarms manually.

Follow these steps:

1. Select the model with the duplicate IP address alarm.
2. Determine whether to allow two devices to have the same IP address. If not, change one of the devices to use a unique IP address, and then use the Update feature to change the IP address within CA Spectrum.

3. To clear the duplicate, click  (Clear selected alarms).

The alarm is cleared. The status color on the model icon returns to a normal green condition unless another alarm is present for the model.

Automatic Naming and Addressing

CA Spectrum implements an automatic model naming and addressing feature through the AUTO_NAME attribute (attribute ID 0x00011979). The value for this attribute is set to TRUE by default for each model type in your modeling catalog. You can disable automatic naming and addressing on a model type basis using the Model Type Editor (MTE) to set the value to FALSE. Otherwise, the feature functions as described below.

If you create a new model using only the IP address, CA Spectrum automatically attempts to supply a name for the model in one of three ways:

- Using NIS (Network Information System) or DNS (Domain Naming Service) to get the name from the modeled device
- Checking the local /etc/hosts file for the name associated with the modeled device's IP address
- Using the IP address as the model name

The priority order of the source that will be used to supply a name for the model (IP Address, Name Service, or sysName) is dictated by the Model_Naming_Options attribute on the VNM model. This attribute can be modified on the VNM model's control view.

Note: For more information about configuring landscapes, see the *Distributed SpectroSERVER Administrator Guide*.

Likewise, if you create a new model and supply only a model name, CA Spectrum will attempt to use NIS, DNS, or the local /etc/hosts file to retrieve the modeled device's IP address.

In either case, as long as the value for `AutoName` is `TRUE` for a particular model type, CA Spectrum will automatically maintain the names for models of that model type as follows: in the event the IP address for a model changes and the original model name was supplied by CA Spectrum, a new name will be supplied using one of the three methods listed previously. However, if the original name was supplied by the user and differs from the name that would be supplied through automatic naming, then the original name will be preserved.

Board and port models are also automatically named by default, each being assigned the name of the parent device model suffixed with the board/port Instance ID. For example, if a model of model type `Hub_CSI_IRM2` is named `IRM2_UK`, and the modeled device has a port with the Instance ID of `2.5`, the name of the port will be `IRM2_UK.2.5`. If the device name is changed to `IRM2_US`, then the name of the port becomes `IRM2_US.2.5`. However, if the device name was user-specified and the user then changes the port name from `IRM2_US.2.5` to `LAB_PORT`, then the automatic naming will not be used for that port in the event of subsequent IP address changes. Some boards (mainly standalone MIMs) contain their own intelligence. In such cases, setting `AUTO_NAME` to `FALSE` as previously described will disable the autonaming intelligence and allow the board's own intelligence to work.


Detection of Firmware Problems

CA Spectrum allows for automatic detection of problems in some device firmware. Using the `connects_to` and `collects` model type relations formed in the Topology views, CA Spectrum detects if a network management firmware problem exists in the hub device. The `connects_to` relation denotes that one model is attached or connected to another; for example, a PC model connects to a hub model via a port on the hub. The `collects` relation denotes that one model collects information from another mode; for example, the Topology view model collects information from the hub devices contained within that view. If CA Spectrum cannot retrieve management information from a hub device, but can still contact devices connected to this hub, then a hub firmware management problem can be deduced and the hub icon's condition color is set to orange. You should then use the alarm views along with the content in the Impact tab and the Performance tab to help isolate and correct the problem.

Interface Intelligence

Interfaces are ports that have both a physical address and a network address. These types of ports are found on routers and bridges, where network identity is important, unlike a port on a repeater, which may have a physical address, but not a network address.

The following diagram illustrates the derivation of the interface port:



Port is the base model type for all ports. Two basic types of ports are derived from Port: repeater ports and interface ports.

Inet_If_Port

Derived from Port and Enet Monitor model types. Inet_If_Port is the base class for all interface ports that are potential monitor points for network statistics.

Data_relay_prt

Derived from Inet_If_Port and SnmpPif. SnmpPif is the base class for all model types that communicate with SNMP agents. Data_relay_prt is the base model type for all interface ports that do their own reading and polling. It is an instantiable model type and is used to model a generic interface port. More specific types of interface ports, such as Cisco_If_Port, are derived from data_relay_prt.

The inference handler CsiHInterfaceIntLinkStatus, which computes InternalPortLinkStatus, is attached at Inet_If_Port so that all interfaces will inherit the desired functionality.

The intelligence for interfaces that use `ifAdminStatus` and `ifOperStatus` are defined in the MIB-II definition of the interface group in RFC 1158. These variables can have the state of ON or OFF, and are defined as follows:

ifAdminStatus: desired interface state

This is the state that the administrator wants the interface to be in. This attribute shows whether or not the interface has been shut off. The values of this attribute are ON and OFF.

ifOperStatus: current interface state

This attribute shows the actual state of the interface. The values of this attribute are UP and DOWN. UP means that the interface is communicating with the network properly. DOWN means the interface has lost connection with the network.

These two variables are used to calculate a CA Spectrum internal attribute named `Internal_Link_Status` (IPLS - 0x10f1b). The possible values for this attribute are `LINK_STATUS_GOOD` (LSG), `LINK_STATUS_BAD` (LSB), and `LINK_STATUS_UNKNOWN` (LSU). This attribute is used to create and clear alarms, both on the interface and the device it is part of. It is also used to generate events concerning the attempt to reach of the interface. The following table shows how these variables are calculated.

IfAdminStatus	IfOperStatus	INTERNAL_PORT_LINK_STATUS
ON	ON	LINK_STATUS_GOOD
ON	OFF	LINK_STATUS_BAD
OFF	ON	LINK_STATUS_UNKNOWN
OFF	OFF	LINK_STATUS_UNKNOWN

The interface's `INTERNAL_PORT_LINK_STATUS` is set to LSU when CA Spectrum has lost contact with the device.

Interface Alarms

Internal Port Link Status (IPLS) is used to generate alarms for both the device model and the interface model. The only interface alarm is gray. These alarms can only be seen by going into the Alarm Details tab of an interface model. It is interesting to note that the alarm for an interface with an IPLS of LSB will have a gray alarm. All models with alarms are displayed in the Alarms tab. This is undesirable for interfaces. Interfaces are considered to be a part of a larger device such as a router. When a router goes down, all of its interfaces goes down as well. A red alarm is generated for the router. It would be confusing to have all of the interfaces producing red alarms as well, cluttering the Alarms tab and making it difficult to locate the router.

A device will watch the IPLS of each of its interfaces. If any of the interfaces has an IPLS of LSB, then the device will generate a yellow alarm with a probable cause of `CS_ALARM_CAUSE_PORT_LINK_STATUS_BAD`. This alarm, once set, will not be reasserted until it is cleared. Only one alarm will be asserted for all ports. The first interface with an IPLS of LSB creates an alarm. The second and successive interfaces with an IPLS of LSB are put into a bad port list. Once the list is clear the yellow alarm is removed.

Each of the interfaces watches its own IPLS. Whenever the interface has an IPLS of LSB or LSU it will generate a gray alarm. This alarm will only show up in the interfaces Alarm Details tab.

When the device becomes unreachable the interface's IPLS is set to LSU. A gray alarm with a probable cause of `CS_ALARM_CAUSE_DEV_CONTACT_STATUS_LOST` is generated.

When the interface has been administratively shut off the interface's IPLS is set to LSU. A gray alarm with a probable cause of `CS_ALARM_CAUSE_ADMIN_SHUT_OFF` is generated.

When the interface becomes unreachable, its IPLS is set to LSB. A gray alarm with a probable cause of `CS_ALARM_CAUSE_PORT_LINK_STATUS_BAD` is generated.

Interface Events

The interface generates two events that deal with its status. The events contain information about the attempts to reach the interface. These events will be generated when a device has a yellow alarm due to a bad interface. Each event message contains the interface number and IP address. For example:

```
Tue 20 Jul, 1994 - 13:31:50 Interface 2 (IP address = 129.128.127.2, type =  
Gen_IF_Port) on device cisc01 of type Rtr_CiscoMIM is unreachable. - (event  
[00010623])
```

As stated before, IPLS has three possible values. This makes it important to know the last two states of an interface's IPLS to make a proper judgment about its current state. If the interface IPLS is LSB, and it was previously LSU, it is important to know if it was previously LSB or LSG. Due to this, each interface keeps the values of the last two states of IPLS. Events are generated based on these two saved values, and the current value.

Events Generated from IPLS State

The following table shows which states generate events based on the IPLS.

Two Previous	Previous	Current	Event
GOOD	UNKNOWN	GOOD	none
GOOD	UNKNOWN	BAD	UNREACHABLE
GOOD	BAD	GOOD	REACHABLE
GOOD	BAD	UNKNOWN	none
BAD	GOOD	BAD	UNREACHABLE
BAD	GOOD	UNKNOWN	none
BAD	UNKNOWN	GOOD	REACHABLE
BAD	UNKNOWN	BAD	none
UNKNOWN	GOOD	BAD	UNREACHABLE
UNKNOWN	GOOD	UNKNOWN	none
UNKNOWN	BAD	GOOD	REACHABLE
UNKNOWN	BAD	UNKNOWN	none

Glossary

Attribute Editor

The *Attribute Editor* is a OneClick utility that lets you change attributes configured at the device level.

configuration

A *configuration* contains the parameters you specify to determine which network entities in your infrastructure you want *Discovery* to locate and identify for review, export, or modeling.

connection

A *connection* is a link between two modeled elements in a view.

container

A *container* is a graphical icon that you can use to depict a group of modeled devices by network technology such as LAN, Network, ATM, or to represent some other containment concept such as a Department.

DCM (Device Communication Manager)

The *DCM (Device Communication Manager)* is the interface between SpectroSERVER and the managed elements. The DCM includes various protocol interfaces that communicate with managed elements using a specific protocol. There is one interface for each of the two supported protocols, SNMP and ICMP. When SpectroSERVER needs to communicate with the managed element, the request is sent on to the appropriate protocol interface in the DCM. The DCM, in turn, passes the request to the managed element.

device attributes

Device attributes are the configuration settings written to a device or interface.

Discovery

Discovery is a OneClick feature that automates the process of discovering and modeling the entities in your IT infrastructure. You can create and edit Discovery and modeling configurations to customize and simplify the process. Discovery also lets you filter and export the results of Discovery or modeling sessions.

Discovery session

A *Discovery session* occurs when you activate a configuration to discover network entities using the parameters specified in the configuration. It uses SNMP and other network technologies to discover and identify network entities specified in the configuration.

Distributed SpectroSERVER (DSS) environment

A *Distributed SpectroSERVER (DSS) environment* consists of more than one *SpectroSERVER*. This environment enables management of a large-scale infrastructure. The SpectroSERVERs in this environment may be located within a single physical location or multiple physical locations. For additional information about a distributed SpectroSERVER environment, see the *Concepts Guide*.

landscape

A *landscape* is the network domain that is managed by a single SpectroSERVER. In OneClick, a landscape is the network view of one SpectroSERVER.

manual modeling

Manual modeling is the act of manually representing individual devices and their connections within a OneClick topology view.

Model by Host Name

Model by Host Name is a modeling feature that you can use in the Universe, World, or TopOrg topology. This feature lets you manually model a new device by specifying the host name for the device.

Model by IP Address

Model by IP Address is a modeling feature that you can use in the Universe, World, or TopOrg topology. This feature lets you manually model a new device by specifying the device IP address.

Model by Model Type

Model by Model Type is a modeling feature you can use in the Universe, World, or TopOrg topology. This feature lets you manually model container icons or devices by a model type.

modeling methods

There are two *modeling methods* available for modeling your network infrastructure. You can automate the process using *Discovery*, or you can manually model the individual entities and later enhance the model presentation using the Topology Edit Mode tools.

modeling session

A *modeling session* occurs when you instruct *Discovery* to model the results of a *Discovery session*. It uses the modeling options specified to model the network entities discovered in that *configuration*.

results list

A *results list* is a detailed list of the network entities discovered from a *Discovery session*, or network entities modeled from a *modeling session*.

security string

A *security string* establishes permission to various elements in OneClick models, such as modeled devices. Administrators can secure access to models using *security strings*.

SpectroSERVER

The *SpectroSERVER* is the server process responsible for providing network management services such as polling, trap management, notification, data collection, fault management, and so on. Also referred to as the *VNM (Virtual Network Machine)*.

topology

A *topology* is an iconic view in OneClick of a modeled network.

VNM (virtual network machine)

An alternate name for the SpectroSERVER, a *VNM (virtual network machine)* is the server process responsible for providing network management services such as polling, trap management, notification, data collection, fault management, and so on.

Index

No index entries found.