

CA Unified Infrastructure Management

Topology and Root Cause Analysis User Guide

v8.0



September 2014

Copyright Notice

This online help system (the "System") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This System may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This System is confidential and proprietary information of CA and protected by the copyright laws of the United States and international treaties. This System may not be disclosed by you or used for any purpose other than as may be permitted in a separate agreement between you and CA governing your use of the CA software to which the System relates (the "CA Software"). Such agreement is not modified in any way by the terms of this notice.

Notwithstanding the foregoing, if you are a licensed user of the CA Software you may make one copy of the System for internal use by you and your employees, provided that all CA copyright notices and legends are affixed to the reproduced copy.

The right to make a copy of the System is limited to the period during which the license for the CA Software remains in full force and effect. Should the license terminate for any reason, it shall be your responsibility to certify in writing to CA that all copies and partial copies of the System have been destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS SYSTEM "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 SYSTEM, 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 manufacturer of this System 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 © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Legal information on third-party and public domain software used in this product is documented in the *Third-Party Licenses and Terms of Use* (http://docs.nimsoft.com/prodhelp/en_US/Library/Legal.html).

Contact CA

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

Send comments or questions about CA Technologies product documentation to nimsoft.techpubs@ca.com.

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

Contents

Chapter 1: About Topology and Root Cause Analysis	7
Overview of Topology and Root Cause Analysis	7
About Network Topology	7
About Root Cause Analysis.....	8
Data Flow Architecture	9
Overview of Set Up and Configuration.....	10
Prerequisites and Supported Platforms	11
Phase 1: Setting up and Verifying Discovery of the Network	11
Phase 2: Initiating Topology Discovery and Verifying the Result	12
Phase 3: Initiating Root Cause Analysis on the Network	12
Chapter 2: Setting up Discovery	13
Setting Device State to "Managed"	13
Chapter 3: Setting up Topology	15
Configuring and Running Relationship Services	15
Configuring and Running the Topology Agent	17
Viewing and Verifying Network Topology	20
Chapter 4: Setting up Root Cause Analysis	23
Installing and Running Topology Fault Correlation	23
Configuring Net Connect	23
Chapter 5: Using Topology and Root Cause Analysis	27
Applying the Network Topology View.....	27
Applying Root Cause Analysis.....	28
Appendix A: Deployment Options	31
Appendix B: Launching Relationship Viewer from the Alarm Console	33
Defining a Launch URL.....	33

Chapter 1: About Topology and Root Cause Analysis

This guide covers high-level concepts surrounding CA Unified Infrastructure Management (UIM) Topology and Root Cause Analysis, and gives an overview of how to put it to work in your environment.

Overview of Topology and Root Cause Analysis

CA Unified Infrastructure Management (UIM) Topology and Root Cause Analysis provides enhanced visibility into the structure and condition of a distributed IP network. It gives you a clear visual representation of your network devices: routers, switches, hosts, hubs, virtual machines, printers, and so on. Moreover, you can see at a glance the condition of those elements and immediately identify devices that are in sub-optimal condition.

About Network Topology

The *topology* of a network describes its elements and their real or virtual connections, irrespective of physical location or other attributes of the network elements. Visualizing the topology of a network makes its organization immediately apparent.

Components in UIM automatically discover and monitor hosts and devices throughout your network, including non-managed devices such as hubs, dumb switches, or devices without SNMP. Using this data, topology deduces the structure of the network and then builds a model of it. This model is ultimately presented to you in the Relationship Viewer portlet within the Unified Monitoring Portal (UMP).

Topology data periodically regenerates to reflect changes on the network, helping track the dynamic changes in a network. Note that what any particular user actually sees in their view is controlled through UIM Server access security.

About Root Cause Analysis

Root cause analysis stems from the non-trivial problem of detecting the cause of a flood of alarms. Often, an otherwise isolated problem triggers a wave of alarms as the effects of the failure ripple across the network. It can be difficult to identify the "root cause" amidst the surge of alarms, most of which are just symptoms of the real issue.

Root cause analysis helps identify the true source of this kind of outbreak. It issues a root cause alarm—which identifies a go-to point for troubleshooting and repairs—and hides the symptomatic alarms that cascade from the initial failure. It also provides visual cues to the affected element(s) in the network topology view.

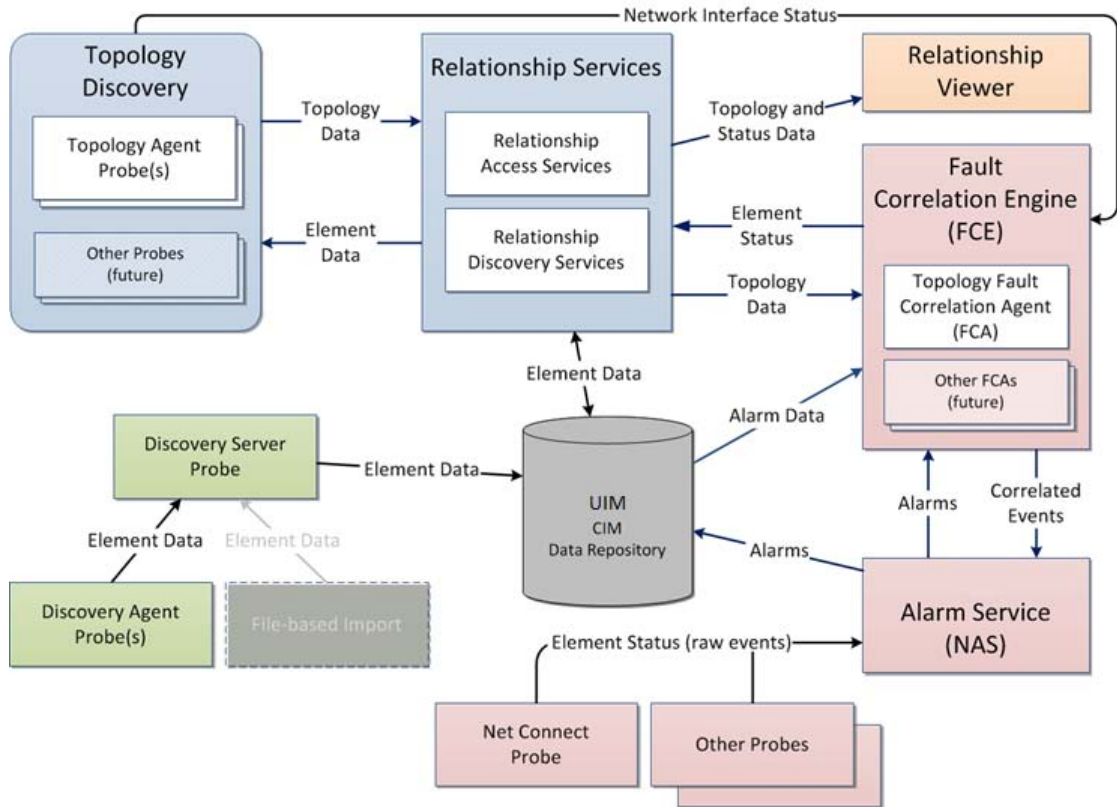
Unlike the network topology, where data is refreshed on a periodic basis, root cause analysis is a continuous process. Elements in the Relationship Viewer in UMP display a status indicator that is a real-time reflection of that analysis.

The status indicators range from "up," which indicates fully normal operation, through increasing levels of sub-optimal status. As alarms flow into the system, root cause analysis determines the actual reason for each, and updates status of affected elements in the network topology view. Whenever conditions in the network or elements change, the status indicators are updated accordingly. One consequence of this is that the effect of repairs is also immediately apparent as Root Cause Analysis resets status symbols appropriately across the network.

For details on using the Relationship Viewer see the Relationship Viewer online help in the Unified Monitoring Portal.

Data Flow Architecture

This diagram illustrates the architectural components that make up Topology and Root Cause Analysis, and how data moves among them.



Note: This diagram does not reflect the actual distribution of probes in the environment; for deployment guidelines see the section on [Deployment Options](#) (see page 31).

Components in the above diagram are either specific to topology and root cause analysis, or in some cases, they are foundational components of UIM which are pivotal to topology and root cause analysis.

The table below describes each component and its role; (color) refers to the color coding in the illustration above.

Service	Probe (or application)	Role
Topology (blue)	topology_agent	This probe formulates network topology using data from the network discovery processes..

	relationship services	This component integrates the raw topology from one or more topology agents, and provides data to the Relationship Viewer.
Root Cause Analysis (red)	fault_correlation_engine	This component provides general fault correlation services. It can host multiple fault correlation applications.
	topology_fault_correlation_agent	This application performs root cause analysis on network state changes using data from the Topology Agent and other sources. Note: The Topology Fault Correlation Agent runs as part of the Fault Correlation Engine, so is not shown separately in the list of probes on the hub robot.
	net_connect	This probe provides node-up and node-down alarms, which feed into root cause analysis.
Discovery (green)	discovery_server	This probe discovers network elements and as much information about these elements as possible.
	discovery_agent	This probe discovers discovery_agents and elements that make up UIM Server. It is responsible for storing information about discovered network elements in the NiS database.
	cm_data_import	This probe is an optional component of discovery, used for file-based device import. Devices imported using cm_data_import are not recognized by topology nor, by extension, root cause analysis.
User interface (orange)	Relationship Viewer	This is the main user interface for Topology and Root Cause Analysis. It presents element and relationship data as well as status information. It runs as a portlet within UMP.

Overview of Set Up and Configuration

Setting up Topology and Root Cause Analysis happens in three major phases, each with several important steps. In some cases you must allow time for the results of one step to be complete, stable, and verified before you proceed to the next.

Note: You use several different tools during the set up and configuration Topology and Root Cause Analysis. This document does not attempt to describe each tool in detail. Instead, it highlights the important steps required in each to accomplish a fully functional deployment of Topology and Root Cause Analysis.

After listing prerequisites, the three major phases of setup are outlined the following sub-sections. Pointers to other documentation and later main sections in this document cover each step in more detail and provide valuable tips to improve performance and simplify your configuration tasks.

Prerequisites and Supported Platforms

Topology and Root Cause Analysis supports the same set of operating systems and databases as supported by the Nimsoft Server solution. Please refer to the [Nimsoft Compatibility Support Matrix](#) for the latest information on supported platforms.

See also the [Support Matrix for Nimsoft Probes](#) for additional specific information on the components that make up topology and root cause analysis.

Viewing network topology requires the Relationship Viewer portlet included in the Unified Monitoring Portal (UMP) 2.1 and above.

Topology and Root Cause Analysis requires that the latest versions of all its various components. Specifically, these packages must all be updated to the latest version available:

- topology_agent
- relationship_services
- fault_correlation
- topology_fault_correlation.

Obtain the latest version of these probes from the [Nimsoft internet archive](#) (<http://support.nimsoft.com>, login required), and confirm they are installed on your primary hub robot.

Note: Be sure to follow the installation and configuration processes described in this document. The ordering of steps is important for a successful implementation of Topology and Root Cause Analysis.

Phase 1: Setting up and Verifying Discovery of the Network

Analysis of network topology requires complete and accurate knowledge of all important devices and hosts in the network. Device discovery is covered in detail in the [Discovery User Guide](#).

Note: File-base import is an option in the discovery process. However, devices added via file-based import are not included in the topology model, nor are they relevant to root-cause analysis.

Phase 2: Initiating Topology Discovery and Verifying the Result

Once you are confident that the network is fully and correctly discovered, you can initiate topology discovery, a process in which the topology of the network is analyzed and modeled. This phase is covered in the section of this guide entitled [Setting Up Topology](#) (see page 15).

Phase 3: Initiating Root Cause Analysis on the Network

With the network correctly and fully represented by the topology visible in the Relationship Viewer, you can initiate root cause analysis. This is covered in detail in [Setting up Root Cause Analysis](#) (see page 23).

Chapter 2: Setting up Discovery

To set the stage for topological analysis, *discovery* of the network, using the discovery server and discovery agent, is the essential first step.

Note: This guide mentions discovery only to describe it as a necessary precursor to Topology and Root Cause Analysis. Device discovery is covered in detail in the [Discovery User Guide](#).

Keep in mind that correctly setting up SNMP is critical for topology:

- A topology agent uses the information gathered from the paired discovery agent to explore the network. It also uses the discovery information to determine whether a node is SNMP capable, and which SNMP authentication profile to use.
- Check carefully to see if all important hosts and network devices are present in discovery, and that all network devices are responding to SNMP queries.

Setting Device State to "Managed"

After you run discovery but before you set up topology, you must change the state of discovered devices to **Managed**. Doing this for each network device or host that you want to include in root cause analysis enables fault correlation for these devices.

As database administrator, execute the following SQL query against your UIM database:

```
update CM_COMPUTER_SYSTEM set state=2 where state=0;
```

Note: State is an attribute that can be set for any system or device. However, only systems and devices that were discovered by discovery agent and discovery server play a role in topology and in root cause analysis.

Chapter 3: Setting up Topology

The next step is to set up *topology*, which consists of configuring relationship services and the topology agent, then verifying that the depicted topology of the network is accurate.

This section contains the following topics:

[Configuring and Running Relationship Services](#) (see page 15)

[Configuring and Running the Topology Agent](#) (see page 17)

[Viewing and Verifying Network Topology](#) (see page 20)

Configuring and Running Relationship Services

The first step in setting up topology is to configure Relationship Services.

Important: To assure optimal performance and reliability, confirm you have the latest version of the Relationship Services component (**relationship_services**) from the UIM internet archive, and that it is installed on your primary hub robot. Note also that the **relationship_services** component uses the "Raw Configure" tool for configuration. It is similar to the MS-Windows registry editor and has no error checking, so take care that any changes you make are valid before you continue.

The default configuration is appropriate for many production environments. However, the following configuration parameters in the Raw Configure tool are of special note:

data-engine

Location: **setup** group

Default: **data_engine**

Description: The default value is appropriate if the data engine probe is running on the same robot as relationship services, which is usually the case. If the data engine probe is running on a different robot, set the **data-engine** configuration parameter to the fully qualified address of the data engine probe:

```
/<UIM_domain_name>/<hub_name>/<robot_name>/data_engine
```

topology-retrieval-interval

Location: **task-config** group

Default: **60000** milliseconds (1 minute)

Description: This specifies how often, in milliseconds (1000ms = 1 second), relationship services checks for topology updates. The default is appropriate for routine operation. During initial set-up you may want to view the new topology soon after an update is available. If so, you can change the **topology-retrieval-interval** to **10000** (10 seconds). However, remember to change it back to the default when you are fully satisfied with everything.

ignore-hosts

Location: **task-config** group

Default: **false**

Description: If set to **true**, topology omits any node that is not a router or switch. This lets you monitor and display only your network backbone, exclusive of any and all host systems connected to it.

allow-unresponsive-hosts

Location: **task-config** group

Default: **true**

Description: If set to **true**, a node that is found to have an active MAC address but which does not respond to ICMP ping or SNMP requests is included in the topology. (Nodes with this behavior may be non-SNMP nodes with firewall restrictions imposed on ICMP requests.) If **false**, a node that does not respond to ICMP ping or SNMP requests is excluded in the topology even if its MAC address is active.

If the **relationship_services** probe is currently deactivated, activate it now.

This completes the configuration of Relationship Services.

Configuring and Running the Topology Agent

The topology agent depends on the discovery agent, even though they don't communicate directly. A topology agent must be co-hosted with each discovery agent in order to use the information gathered from its paired discovery agent.

For example, it uses the discovery information to determine whether a node is SNMP capable, and which SNMP authentication profile to use. This supports *topology discovery*, which is a supplementary scan of SNMP-capable devices to gather more detail on those devices, over and above that found through "regular" phase 1 discovery.

Relationship Services, configured in the previous section, finds the pairs of these agents and provides the information collected by the discovery agent to the corresponding topology agent. The [architectural diagram](#) (see page 9) shows the data flow.

To enable this data flow in your deployment, make sure you have installed a topology agent on every robot that has an actively deployed discovery agent.

Install a Topology Agent Probe

You install a topology agent just as you would install any other probe: using Infrastructure Manager or Admin Console, drag it to the robot where you want a topology agent, and drop it there.

Note: Like any probe, the topology agent is automatically activated upon installation. To adjust the configuration as described below before running it, deactivate the topology agent as soon as it starts.

To Configure a Topology Agent Probe

Note: The `topology_agent` probe uses the "Raw Configure" tool for configuration. It is similar to the MS-Windows registry editor and has no error checking, so take care that any changes you make are valid before you continue.

The default configuration is appropriate for many production environments. However, the following configuration parameters in the Raw Configure tool are of special note:

discovery-interval

Location: `task-config` group

Default: **4320000** milliseconds (12 hours)

Description: This is how often the topology agent runs topology discovery (a deeper query of SNMP-capable devices to gather additional information).

Tip: As a rule, you should use the default interval, or set it to the interval you want used in your final production environment. If you want to run topology discovery more often during your initial set up and testing, it is better to use the `discover-on-startup` parameter rather than the `discovery-interval`.

discover-on-startup

Location: **task-config** group

Default: **false**

Description: Normally the topology agent runs its discovery if it has not already run, or if topology data otherwise does not exist. After that, the topology agent runs on its configured interval and does *not* automatically run on startup.

Tip: If you set this parameter to **true**, you can easily re-run topology discovery at any time--just restart the probe. This is useful when getting started, where you may be re-running the discovery agent and/or topology agent frequently until you get the expected results. When you no longer need to re-run the topology agent, remember to set this option to false.

generate-support-data

Location: **task-config** group

Default: **false**

Description: If set to **true**, the topology agent saves technical support files in the **topology_agent** support directory. This is only necessary if your CA support contact recommends it.

do-interface-polling

Location: **task-config** group

Default: **true**

Description: Enable this option to include network interface status in root cause analysis. This is critical for most users; its recommended setting is **true**.

interface-polling-interval

Location: **task-config** group

Default: **300** seconds (5 minutes)

Description: The frequency with which the topology agent starts a poll of all known interfaces.

Tip: Make a note of the value you enter here, because you will need to know this value when you configure the **net_connect** probe in a later step.

snmp-timeout

Location: **task-config** group

Default: **3000** milliseconds (3 seconds)

Description: The value is in milliseconds (1000ms = 1 second). The default value is generally appropriate unless the network is known to have very high latency.

Note: Devices are dropped from the inventory after the second cycle of topology discovery where they fail to respond within the **snmp-timeout** period.

snmp-retries

Location: **task-config** group

Default: **2**

Description: The default value is generally appropriate. Setting the limit higher might achieve better accuracy in some cases, but at the potential cost of longer topology discovery times. Setting it lower might achieve faster topology discovery times in some cases, but at the potential cost of impaired accuracy.

snmp-walk-enable

Location: **task-config** group

Default: **false**

Description: If set to **true**, the topology agent will perform a full SNMP walk for each IP address and save the results in **walk*.xml** files in the topology agent support directory. The SNMP walk occurs after discovery completes.

Note: Change this parameter only if your CA support contact recommends it, and be sure to reset it to **false** as soon as possible, to avoid unnecessary network overhead.

snmp-walk-max-seconds

Location: **task-config** group

Default: **120**

Description: This parameter applies only if **snmp-walk-enable** is set to **true**. It controls the maximum time that will be spent performing a SNMP walk per IP address. If the SNMP walk doesn't complete within the time limit, the SNMP walk for the IP address is terminated.

If the **topology_agent** probe is currently deactivated, activate it now.

When the above configuration is done on all topology agents in your environment and they are running normally, this procedure is complete.

Viewing and Verifying Network Topology

When topology analysis for the network is finished, you can view the topology using the Relationship Viewer.

Tip: You can learn about using the Relationship Viewer in the UMP help.

As you examine the topology in the Relationship Viewer, look for the following:

1. Are all the expected network devices and hosts included?
 - Start by looking for the main routers and switches on your network. In the relationship viewer, you can search by name or IP address to find nodes. If devices or hosts are missing, try to identify a pattern for nodes that are missing on the map. Do they belong to the same IP subnet or range? If so, you need to adjust the discovery scopes to include the missing nodes.
 - Check for multiple copies of nodes that bridge distinct regions of the actual network. Make sure that no node is discovered by two independent discovery agents.
2. Are network devices and hosts connected as you expect?
 - Are the expected devices or hosts correctly connected to the main routers and switches?
 - Are there islands of nodes that should be connected together?
 - Look at the unconnected nodes: what should they be connected to?
3. Is each network device and host identified by the correct type? For example, are routers identified as routers, switches identified as switches, and so on? It is possible under some circumstance for objects to be incorrectly identified:
 - An actual switch may be misidentified as a host node.
 - An actual switch may be misidentified as an unidentified "device" node.
 - An actual router may be misidentified as a host node.
 - An actual switch/router may be misidentified as switch or router (but not both).
 - An actual host node may be misidentified as a router.
 - A network device may be misidentified as a host node. A discovered object is labeled a host node if it can't be identified as anything else.

Try to identify a pattern for nodes that are incorrectly identified. Are there devices from a particular hardware vendor that are misidentified? Are there particular types of device (e.g. wireless access points) that are misidentified?

If you find incorrectly identified network devices and hosts, check for these possible causes:

- No SNMP information was collected from the device. Without SNMP information, it is not possible to properly identify a discovered node, so it takes on the default identity of host. Potential problems include the following:
 - An incorrect SNMP community string was specified in the Authorization Credentials for the discovery scope for the node's IP address.
 - The SNMP community string was not applied to the discovery scope for the node's IP address.
 - The discovery agent did not complete its SNMP scan of all discovery scopes.
- The device does not allow access to all supported MIBs. For example, if a switch/router is configured to allow access to route table but does not allow access to the bridge MIB, it can only be detected as a router.

When you are satisfied that the topology depicted in the Relationship Viewer accurately reflects the network, you can proceed to the final phase of deployment: activating Root Cause Analysis.

Chapter 4: Setting up Root Cause Analysis

This section describes how to enable *Root Cause Analysis* in your environment.

Begin this process only after you are fully satisfied that the topology represented in the Relationship Viewer is a good match with your actual network. If root cause analysis operates on an flawed model, it can make flawed inferences.

This section contains the following topics:

[Installing and Running Topology Fault Correlation](#) (see page 23)

[Configuring Net Connect](#) (see page 23)

Installing and Running Topology Fault Correlation

The **topology_fault_correlation** application is hosted by the **fault_correlation_engine** probe. The **topology_fault_correlation** application is not a probe and won't be listed by Infrastructure Manager or Admin Console after installing it.

To Install and Run Topology Fault Correlation (if not already installed):

1. In Infrastructure Manager or Admin Console, deactivate the **fault_correlation_engine** service. It should show a gray status indicator.
2. Install the **topology_fault_correlation** application by dragging it to the hub robot and dropping it there.
3. Restart the **fault_correlation_engine** service to activate the new application.

The Topology Fault Correlation application is now running. It is part of the fault correlation engine, so it is not listed among the probes on the robot.

Configuring Net Connect

Fault correlation depends on the **net_connect** probe for node-up and node-down alarm status, so the probe must be configured properly to provide that data.

You can configure multiple nodes at once for network fault correlation using the Unified Service Manager (USM) application in UMP as described below.

Note: There is no need to duplicate `net_connect` monitors. If `net_connect` is *already* configured to ping/monitor one or more nodes, you can skip steps 1 through 7 below for *those* nodes. You must configure the **Check Interval** and **interface-polling-interval** for `net_connect` against those nodes, however. Use the interval settings described in step 3e below.

To Configure Net Connect for Network Fault Correlation

1. In USM, create a new group that includes all of the IP addresses to be monitored. If all IP addresses are part of the same hub or origin, you can define a filter using the origin to simplify creating the group, and dynamically maintain its membership.
2. Enable the **Edit Monitoring Templates** tab in USM by executing the following three SQL queries against your NiS database.

The examples below are for generalized MySQL. MS-SQL and Oracle syntax may vary.

- a. Execute this query:

```
select distinct template_id from cfg_template
```

Make a note of the `template_id` for the top row returned.

Note: If no rows are returned, re-deploy the ACE probe using Infrastructure Manager or Admin Console, then execute the query again.

- b. Execute this query:

```
select distinct grp_id from cm_group
```

Identify the group you created in step one from the list that is returned. Make a note of its `grp_id`.

- c. Execute this query:

```
insert into cfg_group_template (grp_id, template_id,  
template_order) values (<value from step 2>, <value from step  
1>, 0)
```

- d. Restart USM (from the wasp GUI).

- e. Reload USM in your browser. You should now see the **Edit Monitoring Templates** tab available for use.

3. Create a monitoring template as follows:

- a. Close **MY NETWORK** in the upper left of the USM and open **ADMINISTRATION**. (Clicking on the pencil icon is the equivalent).
- b. Add a new monitoring template (click the plus sign icon) and name it as desired.
- c. Click the **Add** button to add a monitor.
- d. Select the **Ping (remote)** monitor in the **net_connect** package and click **OK**.
- e. With the monitor selected in the main USM window, click **Edit** to configure its settings as follows:
 - Set the **Check Interval** for the monitor to *double* the length of the **interface-polling-interval** (see page 17) set for the topology agent. This assures the responsiveness of topology fault correlation in performing its root cause analysis. CA recommends setting Check Interval to 600 seconds, and the interface-polling-interval to 300 seconds.

- Use the check box to enable **Monitor ICMP connectivity (ping)**. Within this box, set the parameters as follows:
 - Set the **Packet Size** to **default**.
 - Set **Retries in interval** to a value of **1**.
 - Set the remaining configuration parameters in the **Monitor ICMP connectivity (ping)** box as appropriate for your situation.
 - Use the check box to enable **Generate alarm**. Within this box, set the parameters as follows:
 - On the **Connectivity** tab, set the **Timeout (sec)** value to **1**, set the **Failed Intervals** to **1**, and set the **Identification Method** to **IP Address**.
 - Set parameters on the **Threshold** and **Packet Loss** tabs as desired.
4. Remove the template object that was created in step 2 from the group you created, leaving the monitoring template you added in step 3.
 - a. Select the group you created in the left frame of USM, then click on the gear icon to the right.
 - b. Click on **Monitoring Templates**.
 - c. Move the template object created in step 2 to the right list "Available," and conversely move the template you created in step 3 to the left "Linked Monitoring Template."
 - d. Click **Okay**.
 5. Close **ADMINISTRATION** and open **MY NETWORK** again.
 6. Select the group you created in step 1.
 7. Apply the monitoring template to the group by dragging it to the group and dropping it there.

It may take several minutes for the above changes to be reflected in the **net_connect** probe. To verify that your settings have been applied, use the Infrastructure Manager or Admin Console to view the status of the **net_connect** probe.
 8. Ensure that all nodes in the group have their state set to **Managed**. See the section Setting Device State to Managed for more information.

This completes the deployment of Topology and Root Cause Analysis. All components should be fully and correctly configured at this point.

If you made temporary configuration changes during the deployment, you should reset them to values that are appropriate in a production environment. In particular, you may want check the following:

- On each discovery agent in the system, review the effectiveness of discovery using non-SNMP protocols. Ineffective protocol requests slow down discovery. Depending on the time spent for a given protocol and its effectiveness, you may want to adjust its settings, or consider disabling it.
- Set the [topology-retrieval-interval](#) (see page 15) to its default value, or a value you prefer.
- Check the [discover-on-startup](#) (see page 17) parameter of each topology agent; in a production environment, it should be set to **false**.

Chapter 5: Using Topology and Root Cause Analysis

This section describes how to apply Topology and Root Cause Analysis to address issues in your IT environment.

Applying the Network Topology View

With Topology and Root Cause Analysis fully functional, you can use the Relationship Viewer to view the topology of the network in various ways. You should explore the features of the Relationship Viewer to see how they apply in your situation.

For example, suppose you have set up a group in the Unified Service Manager that contains all the connector devices on the network backbone (switches, routers, hubs, and so on), and which excludes everything else (workstations, network printers, servers, and so on).

If you select that group in USM and then launch the Relationship Viewer, the display shows the central structure of the network. By adjusting the **Radius** slider in the Relationship Viewer, you can easily discover the key connections that are pivotal to your entire network.

Similarly, you might have another group in USM containing the servers that provide web presence to the internet. By choosing that group, and incrementally expanding the radius of the view, you can similarly visualize how those servers reach the internet cloud.

Explore the different layouts in the Relationship Viewer, which provide different insights into the structure of the network. Note, however, that some layouts are more suited to IP networks than others; some exist primarily for use in future applications.

Applying Root Cause Analysis

The Relationship Viewer uses symbols to represent network elements and various other devices. Each symbol includes an indication of the status of the associated object. The status indicator for an element reflects the most severe current condition known for that object. There are two kinds of status:








Quality of Service Status



Quality of Service (QoS) status indicators signify that at least one QoS alarm has been issued for the object. The device is otherwise functioning normally on the network.

Operational Status

Operational status indicators reflect abnormal network behavior of the element that render any QoS alarms for it temporarily irrelevant.

The following table presents the status indicators in their order of severity.

Status Indicator	Name	Description
	Up	There are <i>no</i> QoS alarms for the device and it is functioning normally.
	Informational	There is at least one informational QoS alarm for the device.
	Minor	There is at least one minor QoS alarm for the device; there may also be QoS alarms of lower severity.
	Major	There is at least one major QoS alarm for the device; there may also be QoS alarms of lower severity.
	Critical	There is at least one critical QoS alarm for the device; there may also be QoS alarms of lower severity.
	Administratively Down	The device is temporarily out of service for maintenance.
	Unreachable	The device can not be reached. Its current condition is unknown, but it has not been determined to be at fault. If and when Root Cause Analysis determines it actually <i>is</i> at fault, the status indicator changes to the Fault symbol. Otherwise, the status indicator changes only when the device is again reachable and its status at that time can be determined.

Status Indicator	Name	Description
	Fault	The device has experienced a serious malfunction that can probably only be resolved with intervention.
	Unmanaged	The device is unmanaged, and could be in any state. An unmanaged device is not monitored, so its actual condition is never ascertained.

When a network device (router, switch, etc.) fails, its status indicator is initially set to "Unreachable". It remains so until Root Cause Analysis determines that the device is the source of the problem and sets its status indicator to "Fault."

Root Cause Analysis also sets the status of devices that are beyond the fault—from the point of view of the robot performing analysis—to "Unreachable." As soon as the connector device returns to normal operation, Root Cause Analysis recalculates the status of the "Fault" and "Unreachable" nodes based on their condition at that time and sets status indicators appropriately.

The status indicators displayed in the Relationship Viewer help you quickly evaluate the repercussions of a given failure. Knowing the effect of a particular failure gives you immediate insight into its priority for attention.

Appendix A: Deployment Options

The components that constitute topology and root cause analysis are included in a standard installation of CA Unified Infrastructure Management (UIM) Server. All these components are initially located on the UIM Server--some can be replicated and deployed to remote machines as needed.

Probe (or application)	Role	Deployment Options
discovery_agent	Handles network discovery	Service providers and those with very large networks may find it useful to deploy multiple discovery and topology agents in various locations. This will divide discovery of a large network across administrative boundaries—so that different users have access to different parts of the network—or in situations where there is no direct connectivity to devices at a remote site because of firewall constraints or network-address translation (NAT).
topology_agent	Performs topology analysis	Can be distributed as multiple instances. Note that every running topology agent <i>must</i> be paired with a discovery agent on the same robot.
relationship services	Handles communication between discovery_agent and topology_agent; connects many topology and RCA processes	Recommended to be kept together with the fault_correlation_engine on the UIM Server.
fault_correlation engine	Performs root cause analysis	Recommended to be kept together with relationship_services on the UIM Server.
topology_fault correlation	Performs network fault analysis	Hosted by the fault_correlation_engine probe on the UIM Server.
discovery_server	Persists discovery data in the NiS database	Must be deployed on the primary hub robot of the UIM Server.
net_connect	Provides node-up and node-down alarms	Deployed in the background by the ACE probe as required (following configuration for RCA); covered in the section Configuring Net Connect (see page 23).
Relationship Viewer	User interface in UMP for viewing topology and RCA	Recommended deployment for UMP is on a host separate from the UIM Server.

Appendix B: Launching Relationship Viewer from the Alarm Console

This section describes how to configure the Alarm Console so that you can launch the Relationship Viewer from alarms, with the associated element centered in the network topology view.

Defining a Launch URL

This section supplements information found in the online help for Relationship Viewer, and the online help for the **New Action** dialog of the Alarm Console. Look at the information in those locations for details about features mentioned only briefly here.

The documentation for the Relationship Viewer explains how to build a URL to launch the Relationship Viewer. The URL parameters described there have been extended as follows:

- The **type** parameter also allows the value **alarmid**. For example:
`...type=alarmid...`
- If the type specified is **alarmid**, as in the example above, the **elements** parameter specifies a replaceable parameter (**\$ID**) that gets replaced at launch by the alarm identifier. The Relationship Viewer resolves the alarms ID to the individual device or host in the topology, and displays it centered in the view. See the example URL below.

Tip: Specify a radius of **1**, so that the view is restricted to the specified systems.

Use the **New Action** dialog of the Alarm Console window to create a new action that uses the URL, as described in the **New Action** dialog online help. Then proceed to create the new action as usual.

When you are done, you can thereafter select an alarm in the Alarm Console, right-click to pop up a menu of actions, and from that launch the Relationship Viewer with the target system centered.

Example

The following URL is an example of a correctly formed URL, where **<host>** indicates the name of the system running the Unified Monitoring Portal:

```
http://<host>/relationshipviewer/jsp/standalone.jsp?type=alarmid&elements=$ID&radius=1&relationship=physical_connection&sid=$SESSIONID
```

Use the above example to create the value of the **URL** field of the **New Action** dialog of the Alarm Console. Enter the replaceable parameters (**\$ID** and **\$SESSIONID**) exactly as shown, or select them from the list of replaceable parameters in the **New Action** dialog.

Glossary

discovery

Discovery is the process of identifying devices within an IT environment through the use of communication protocol pings and queries. Specific to UIM, discovery is the automated discovery of hosts and devices throughout a network, recording any device within a discovery scope that responds to a request on any configured protocol, including a simple ICMP ping.

UIM Discovery Wizard allows you to set authentication credentials and define IP address scopes to scan on-demand, or be scheduled to run on regular intervals. Protocols used are ICMP, ARP, DNS, SNMP (v1, v2 and v3), WMI, SSH, and NetBIOS.

root cause analysis

Root cause analysis is the process of identifying the single event, such as the failure of a network device, that precipitates a multitude of alarms. By hiding those secondary alarms and issuing a *root cause* alarm for the failed device, Root Cause Analysis focuses attention on the true problem.

topology

The topology of a network is a depiction, usually graphical, of its elements and their real or virtual connections, irrespective of the physical location or other attributes of the network elements.

topology discovery

A supplementary discovery scan, performed by the `topology_agent`, of SNMP-capable devices to gather more detail on a deeper level than that found through "regular" phase 1 discovery.