Cisco Configuration Assurance Solution Audit and Analysis IT Sentinel User Guide

Software Release 11.0.A

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: OL-7550-01

Cisco Configuration Assurance Solution
Audit and Analysis
IT Sentinel User Guide

Copyright

Document Copyright

Title: User Guide
Part Number: D00265
Version: 5

Software Copyright

Product Name: IT Sentinel
Product Release: 11.0

Documentation Conventions

OPNET documentation uses specific formatting and typographic conventions to present the following types of information:

- Objects, examples, and system I/O
- Object hierarchies, notes, and warnings
- Computer commands
- Lists and procedures

Objects, Examples, and System I/O

- Directory paths and file names are in plain Courier typeface:

  `opnet\release\models\std\ip`

- Function names in body text are in italics:

  op_dist_outcome()

- The names of functions of interest in example code are in bolded Courier typeface:

```
/* determine the object ID of packet's creation module */
src_mod_objid = op_pk_creation_mod_get (pkptr);
```

- Variables are enclosed in angle brackets (< >):

  `<opnet_user_home>/op_admin/err_log`

Object Hierarchies, Notes, and Warnings

Menu hierarchies are indicated by right angle brackets (>); for example:

Open File > Print Setup > Properties...

Attribute hierarchies are represented by angled arrows (↘) that indicate that you must drill down to a lower level of the hierarchy:

Attribute level 1 ➘ Attribute level 2 ➘ Attribute level 3

Note—Notes are indicated by text with the word Note at the beginning of the paragraph. Notes advise you of important supplementary information.

WARNING—Warnings are indicated by text with the word WARNING at the beginning of the paragraph. Warnings advise you of vital information about an operation or system behavior.

Computer Commands

These conventions apply to Windows systems and navigation methods that use the standard graphical-user-interface (GUI) terminology such as click, drag, and dialog box.

- Key combinations appear in the form "press **<button>+x**"; this means press the **<button>** and **x** keys *at the same time* to do the operation.
- The mouse operations *left-click* (or *click*) and *right-click* indicate that you should press the left mouse button or right mouse button, respectively.

Lists and Procedures

Information is often itemized in bulleted (unordered) or numbered (ordered) lists:

- In bulleted lists, the sequence of items is not important.
- In numbered lists, the sequence of items is important.

Procedures are contained within procedure headings and footings that indicate the start and end of the procedure. Each step of a procedure is numbered to indicate the sequence in which you should do the steps. A step may be followed by a description of the results of that step; such descriptions are preceded by an arrow.

Procedure FM-1 Sample Procedure Format

1 Procedure step.

➥ Result of the procedure step.

2 Procedure step.

End of Procedure FM-1

For more information about using and maintaining OPNET documentation, see the OPNET IT Sentinel Documentation Guide.

Document Revision History

Release Date	Product Version	Chapter	Description of Change
February 2005	11.0 PL3	Importing Topologies	• VNE Server Import modified for changes to collected configuration import.
October 2004	11.0 PL1	Viewing Networks	• Updated Network Layout and Appearance and View Menu information, including new Layout Nodes Automatically (Simple) command.
		Importing Topologies	• Updated DCI and VNES Import sections. • VNE Server Import significantly enhanced.
August 2004	11.0	Customizing Models	• Updated information on setting attribute values, including new sections Using the Attribute Template Method and Defining an Attribute Template.
		Modeling Traffic	• Added mention of Point Demands.
		Publishing Results	• Updated information on User-Defined Reports, including new procedure for live reports.
		Importing Topologies	• Updated DCI and VNESI information for 11.0. Most of chapter was replaced.
January 2003	10.5	All	First issue

Contents

List of Figures

List of Tables

List of Procedures

1 Introduction to IT Sentinel

IT Sentinel is a comprehensive product suite that can monitor, analyze, and troubleshoot your network automatically. IT Sentinel consists of tools that can identify recent changes in your network, validate network behavior, identify violations of configuration management policies, ensure network integrity, and identify critical errors before they result in undesirable behavior. IT Sentinel can also generate detailed reports that show network changes, device configurations and misconfigurations, security gaps, and other critical information about your network.

Product Features

Key features of IT Sentinel include:

- VNE Server—VNE Server collects data from a variety of network-monitoring tools, removes outdated information, and maintains a unified view of your production network. IT Sentinel can import topology and traffic data from VNE Server to create an accurate and detailed model of your network.
- Device configuration import—An alternative to importing from VNE Server is to import configuration data from Juniper routers, Cisco routers, and Cisco Catalyst switches. IT Sentinel uses this data to create a network topology that models the routing and switching behavior in your production network.
- NetDoctor validation—NetDoctor employs a comprehensive “rules engine” that validates configurations throughout your network and checks for misconfigurations, protocol conflicts, and other violations of good practice.
- Flow Analysis—The IT Sentinel Flow Analysis engine calculates the
  - — Routes taken by each circuit and traffic flow in your network
  - — Resulting link utilizations

  Flow Analysis identifies unroutable flows and overutilized elements automatically.
- Automated reporting—IT Sentinel can generate several different types of reports:
  - — Reports of NetDoctor and Flow Analysis results
  - — User-defined reports that show information relevant to specific network or device configurations
  - — Network difference reports that compare network models imported over time and identify differences in topology and configuration

  IT Sentinel stores all reports on a centralized server, which users can access using a web browser.

Typical Applications

You can use IT Sentinel to do a variety of network monitoring and troubleshooting tasks. The following list describes some possible applications and the product features that support them.

- Track network changes—IT Sentinel can save the current network state as reflected in each individual run. For example, you might want to schedule a specific task to import your network from VNE Server every night. As part of the task, you can generate a network difference report that shows all devices and configurations that have changed since the previous night's run.
- Enforce network configuration policies—You can schedule IT Sentinel to do systematic audits of your entire network using NetDoctor's rules-based engine. NetDoctor has an extensive rules suite that addresses problems in several areas: IP addressing and routing, protocol configurations, IP QoS, SNMP, and so on. You can even create and apply your own custom rules to codify local configuration policies and best practices.
- Audit network security—NetDoctor can find security errors that prevent valid connectivity or fail to block unauthorized traffic. NetDoctor can also verify that the packet filters and route maps referenced by interfaces are defined correctly.
- Find network bottlenecks—If a task imports traffic data from VNE Server, IT Sentinel can run Flow Analysis to identify the bottlenecks at current traffic levels.

Workflow Description

The basic IT Sentinel workflow is to set up a sequence of operations (called an *automation task*), then schedule the task to run automatically at specified intervals. A typical task includes the following operations:

1) Import data from configuration files or VNE Server and create a model of the current network.

   IT Sentinel generates data by analyzing a model of your production network. This model characterizes essential information about your network's topology, configurations, and (in some cases) traffic. Usually the first step in an automation task is to create a new model that reflects your current network.

2) Analyze the network to find overutilized links, misconfigured nodes, and other problems or potential problems.

   After Sentinel creates a network model, it can do a variety of analyses of your network's configuration and performance. For examples of analyses you might include in an automation task, see Typical Applications on page ISU-1-2.

3) Generate reports of the analysis results and store the results on a centralized report server.

   An automation task usually exports results to one or more web reports. IT Sentinel can generate a variety of reports to show network changes, protocol misconfigurations, security gaps, and other data of interest.

IT Sentinel's automation capabilities enable you to set up a task once and schedule it to run repeatedly. Depending on your needs, you might want to define multiple tasks so that each task performs a specific function and runs on a unique schedule. After you define and schedule a task, there is no further need to interact directly with the IT Sentinel software. Any authorized user can see Sentinel results simply by logging on to the report server.

References

Table 1-1 lists references to relevant sections in the IT Sentinel documentation.

Table 1-1 IT Sentinel: References (Part 1 of 2)

Topic	Reference
Feature overviews:	
• Automation Module	Chapter 1 Overview on page AUG-1-1 of the Automation *User Guide*
• VNE Server	VNE Server *User Guide*
• NetDoctor	Chapter 1 Overview on page ND-1-1 of the NetDoctor *User Guide*
• Flow Analysis	Chapter 1 Overview on page FA-1-1 of the Flow Analysis *User Guide*
• Report Server	Report Server User Guide
Automation operations:	
• Automation workflow	Workflow Description on page AUG-1-1 of the Automation *User Guide*
• Automated operations	Creating Task Steps on page AUG-2-1 of the Automation *User Guide*
• Automation tasks	Creating and Scheduling an Automation Task on page AUG-2-5 of the Automation *User Guide*
• Viewing the results of automation runs	Automation Logs on page AUG-2-13 of the Automation *User Guide*

Table 1-1 IT Sentinel: References (Part 2 of 2)

Topic	Reference
Network imports:	
• VNE Server imports	VNE Server Import on page ISU-10-2
• Device configuration imports	Device Configuration Imports (DCI) on page ISU-10-15
Network analysis operations:	
• Configure/Run NetDoctor	Configuring NetDoctor on page ND-3-4 of the NetDoctor *User Guide* Generating a Report From a Template on page ND-3-11 of the NetDoctor *User Guide*
• Configure/Run IP Flow Analysis	Configuring Flow Analysis on page FA-2-4 of the Flow Analysis *User Guide*
Network reports:	
• Generate network-difference report	Network Difference Reports on page ISU-11-5 of the Sentinel *User Guide*
• Generate user-defined report	User-Defined Reports on page ISU-11-3 of the Sentinel *User Guide*
• Flow Analysis reports	Selecting Reports on page FA-2-9 of the Flow Analysis *User Guide*
• Failure Analysis report	Configuring and Running a Failure Impact Analysis on page FA-3-1 of the Flow Analysis *User Guide*
End of Table 1-1	

2 Getting Started

With IT Sentinel, you can create and study computer models of your network. A network model characterizes essential information about the network topology, network traffic, and the network configuration.
You can create a model by

- Importing data from systems that monitor the topology and traffic of your network
- Importing configuration data from device configuration files

After you create a model, you can

- Perform a variety of analyses to validate network behavior
- Identify misconfigurations and errors
- Identify performance problems

Modeling a Network Topology and Network Traffic

This section contains the following sections:

- Creating Network Topologies
- OPNET Model Library

Creating Network Topologies

The Project Editor window provides an environment for defining the topology of a communication network. A network model consists of *nodes* that communicate with each other through *links*.

With IT Sentinel, you can import complete topologies using device configuration files and OPNET VNE Server import. Table 2-1 lists topics in the document set that apply to importing topologies and where you can find that information.

Table 2-1 Creating Topologies

Topic	Reference
Network models, subnetworks, nodes, and links	Chapter 5 Network Overview on page ISU-5-1
Working with network objects	Creating and Editing Network Objects on page ISU-6-8
Importing topologies	Chapter 10 Importing Network Data on page ISU-10-1
Strategies and methodology	Methodology*: Creating a Unified View of the Network*
End of Table 2-1	

OPNET Model Library

Every network object is based on an underlying model.

With IT Sentinel's extensive model library, you can represent any kind of equipment in a communications network: workstations, servers, hubs, routers, switches, bridges, and other network devices. The OPNET model library also includes a comprehensive set of protocol models for studying the effects of different protocol configurations.

Workflows for Common Use Cases

The following table lists several types of studies that you can do using IT Sentinel. A description and suggested workflow for each of these studies, called use cases, follows the table. Use the workflows as a starting point when beginning a new study. The workflows show the steps you can use to create and analyze a use case with IT Sentinel and its add-on modules. The modules used for each use case are also listed in the table. If you do not have a required module, you can modify the workflow to include only the modules you have.

Table 2-2 Workflows for Common Use Cases

Use Case	Objectives
Failure Impact Analysis on page ISU-2-4	Determine the disruptions caused by particular node and link failures. Determine which failures are not 100 percent survivable.
Unified View of Network Information and Intelligent Reporting on page ISU-2-5	Visualize and report on network topology, configuration, and traffic in one environment. Provide unified inputs for other intelligent network management use-cases.
End of Table 2-2	

Failure Impact Analysis

User Groups in Typical Organizations:

- Network operations
- Server operations

Failure impact analysis enables organizations to understand how network performance is affected due to planned and unplanned outages. This consists of running exhaustive failure analyses to determine which traffic flows will be most affected by outages, and where the resulting bottlenecks might be.

This use case examines how the network reacts to node and link failures in the network. You can determine:

- if the network will meet performance objectives under failure conditions
- if the network is vulnerable to certain failures
- which traffic flows are vulnerable to failure

Procedure 2-1 Failure Impact Analysis

1 Configure Flow Analysis—Configure the flow analysis runs that occur during the failure impact analysis.

For more information, see Configuring and Running a Flow Analysis on page FA-2-4 in the Flow Analysis *User Guide*.

2 Configure the Failure Impact Analysis—Configure the failure scenarios you want to analyze by configuring which objects fail and whether they are failed 1-by-1, in pairs, or all at once.

For more information, see Configuring and Running a Failure Impact Analysis on page FA-3-1 in the Flow Analysis *User Guide*.

3 Run the Failure Impact Analysis.

4 Analyze results—Look at the reports from the failure impact analysis to see if the network meets performance objectives during failure scenarios. For information about viewing and using reports, see the Report Server User Guide.

End of Procedure 2-1

Unified View of Network Information and Intelligent Reporting

User Groups in Typical Organizations:

- Network operations
- End-to-end performance studies

Operators of large enterprise and service provider networks have an overwhelming amount of network data, but they rarely have an up-to-date, accurate picture of the network that they can rely on. Having an on-line, continuously valid, integrated view of your network plays a critical role in enabling advanced problem resolution, problem prevention, and network planning. This is possible by collecting network data from different—often disparate—sources, and intelligently merging this information to create a unified network representation for network planning, network engineering, and operations.

3 The Model Library

This chapter describes the features of the standard model library. The results of the simulations you run in OPNET depend on the models you use and the parameters you assign to the models.

To help you build models that achieve their purpose, this chapter contains

- General concepts applied throughout the OPNET standard model library—the library of models available to you as “plug-and-play” objects
- Tips and techniques that enhance your productivity when working with many of the models in the library

You can use models from many sources in your IT Sentinel simulation projects:

- Standard Models—OPNET maintains an extensive library of models that you can use to build networks. These models, which are included on the release CD, are called “standard models”.

  Standard models are defined with a set of functions that are typical in many network environments. You can also develop your own models and share them with other OPNET users.

Organization of the Standard Model Library

The standard model library consists of the following types of objects:

- Devices
- Links
- LANs and clouds
- Demands
- Paths

The library also contains many models of networking protocols and algorithms that allow your network models to simulate real network behavior. However, as a user of OPNET, you do not manipulate the internals of these protocols directly. Instead, you have access to protocol functions through parameters, much as you do in working with real networks. Their parameters appear as attributes of the objects and are configurable via OPNET's GUI. One of the goals of this chapter is to explain the objects you have available to you and how to work with their attributes to represent your networks as desired within the tool.

Device Models

Devices comprise the majority of the objects in the standard model library. They correspond to a wide class of network hardware, including the following:

- Routers
- Switches
- Hubs
- Workstations
- Servers
- Firewalls
- Printers

As you can see, devices essentially correspond to the "boxes", that is, the "chassis-type" or rack-mounted systems in your network. They represent the hardware that performs the information transmission and processing in the network, ranging from simple repeaters, like hubs, to content and computation servers, like mainframe computers. Of course, these devices must not consist solely of hardware; most of them contain large amounts of software models, spanning appropriate layers of the protocol and application stack. What's inside a device depends on its function. For example, a typical router model will contain hardware and software models of Ethernet, PPP, and perhaps other Layer 2 protocols. It will also contain routing protocols such as RIP, OSPF, IGRP, and BGP4.

The standard model library presents device models to you graphically. Typically, you will select these objects from the object palette where they appear graphically, though in some cases you may choose their name from a menu of available models. After they are deployed in your network model, devices appear as icons. The standard model library follows conventions (third party or supplemental models may not) for the appearance of each type of device, as illustrated below:

Figure 3-1 Graphical Conventions for Network Objects

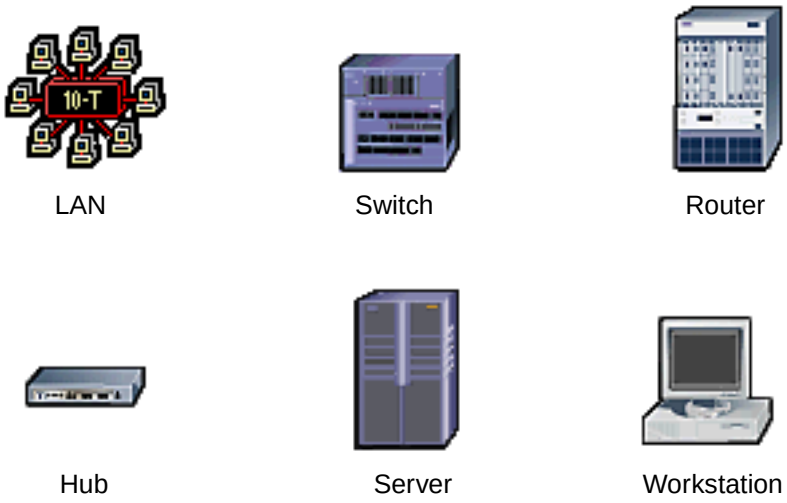

Device models are also categorized into two classes:

- vendor device models
- generic device models

Vendor Device Models

Vendor models represent devices manufactured by a particular company, such as Cisco Systems or 3Com. The developers of the standard model library use data published by these manufacturers to characterize the devices as well as they can. You can also create vendor device models using the “Device Creator” operation, provided you have obtained values for required parameters of the device.

Generic Device Models

Generic models provide behavior that is appropriate for devices of their class. However, they are not configured to model any particular manufacturer’s devices. Instead these devices provide attributes (that is, parameters), allowing you to configure each one you deploy differently if you choose. For example, the generic router below offers the attribute “forwarding rate”, which specifies the throughput of the router in packets per second. Each instance of the router that is deployed can be assigned its own forwarding rate. In contrast, a vendor

device model of a router would already be aware of its own forwarding rate, because the device type is already known. A vendor model would therefore have a preconfigured “forwarding rate” attribute that is consistent with the actual router’s forwarding rate.

Figure 3-2 Generic Model vs. Vendor Model

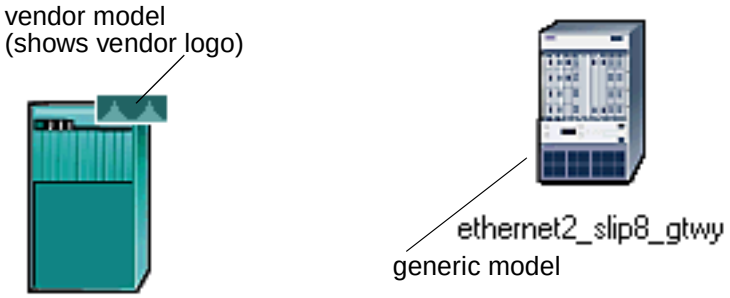

In later sections of this chapter, we will discuss practical aspects of working with the devices in the model library.

Link Models

To form a network of devices, you will need to use links, and these links will require specific characteristics. In OPNET links represent the physical media and properties, such as line rate in bits per second, delay, and likelihood of data corruption. Link models also generally represent a choice of Layer 2 technology, allowing OPNET to verify compatibility of two or more attached devices and the link that connects them. One of the most important characteristics of a link model, from a network performance perspective, is the speed of transmission in bits per second. This characteristic is usually implicit in the choice of link model (for example, a 10BaseT link automatically provides for a 10 Mb/sec. transmission rate).

Links are represented as line segments or a series of line segments with arrowheads in the OPNET GUI. When selecting links from OPNET’s object palette, you will see objects similar to those shown below:

Figure 3-3 Link Model Graphical Convention

LAN and Cloud Models

OPNET lets you model the “end systems” of your network in detail, representing each device if necessary. However, in many simulation studies it is adequate to abstract local area network infrastructure into one object, called a LAN object. The LAN object models many users and a server on the same LAN. However, it does so with one object, which dramatically reduces the amount of

configuration you need to do to represent your internetwork of LANs. In addition, because fewer objects are present in your network, LAN objects help reduce the amount of memory needed to run simulations. LAN objects are provided for a variety of local area network technologies, as shown below:

Figure 3-4 LAN Objects Abstract Local Area Infrastructure

Using LAN objects you can quickly generate large amounts of users for your network. Specifically, each LAN object allows you to specify the number of users present within it. You can then assign application traffic to a subset, or all of the users of the LAN. Thus, scaling the traffic generated by a LAN (to model more users, for example), is a simple matter of increasing the specified number of users. LANs can then be interconnected via switches and routers to carry traffic to and from devices and LANs in other parts of the network.

In a manner similar to the use of LANs, it is sometimes appropriate to abstract parts of the wide area network infrastructure. Cloud models are special objects in the model library used to represent such infrastructure. They provide high-level characteristics used to simulate the behavior of that portion of the network. The ATM, Frame Relay, and IP model suites all include cloud models.

Figure 3-5 Cloud Objects Abstract WAN Infrastructure

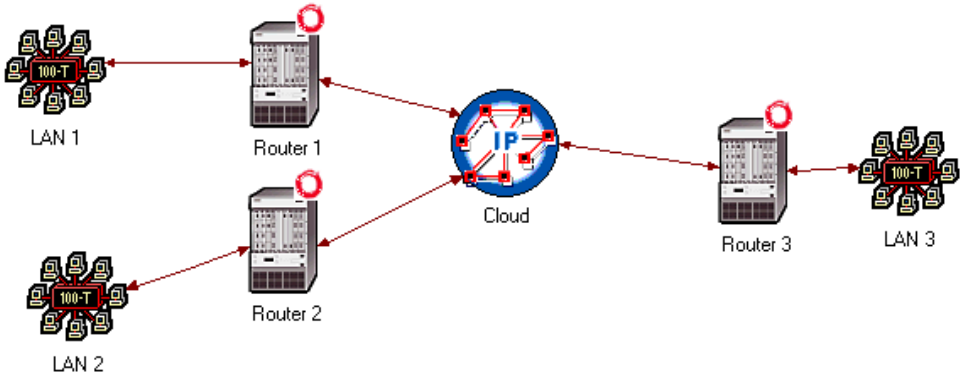

Cloud models can have numerous applications. For example, your backbone network may depend on services delivered by an Internet Service Provider (ISP), which in turn may be connected to a carrier network. Because you do not know the details of the backbone, a cloud node gives you a simpler model, without losing any detail.

You can simulate a backbone network using two cloud-model attributes:

- Packet (or Cell) Latency—This specifies the one-way delay that each packet experiences while traversing the network. You can run a simple test on your network to determine this setting. First, configure your ping traffic (using an option like IP's "record route") to report the number of hops that a packet traverses. Then measure the response time and number of hops for a typical ping packet sent across the network. (Keep in mind that this response time includes transmission delays at each hop.) When you have these two values, you can use the following equation to compute the actual network latency:

  Latency = (Ping Response Time – (Hop Count * (Ping Packet Size/Data Rate)))/2

  — Use the resulting value for the latency parameter. You can also specify any built-in or custom distribution to model variability in the delay.

- Packet (or Cell) Discard Ratio—This specifies the ratio of packets dropped to packets submitted to the network backbone. You can set this value based on your network service provider's statistics or its traffic-contract guarantees.

General Advice When Working with the Model Library

This section presents some general information about working with models in the library. Most of the features and techniques described apply to all of the models.

Table 3-1 Standard Model Library Protocols and Technologies

Layer 1, 2 and Support	Layer 3 and Support	Layer 4
		TCP
Ethernet 10, 100, 1000	IPX	UDP
ARP	Frame Relay	NCP
Frame Relay	IP	
FDDI	RSVP	
Token Ring 4, 16	OSPF	
PPP	RIP	
SLIP	BGP4	
SRP	IGRP	
Spanning Tree	EIGRP	
	IS-IS	
End of Table 3-1		

4 User Interface

This chapter describes the IT Sentinel user interface (UI) and how to use it. The IT Sentinel user interface fits smoothly into the context of common window managers by providing conventional objects such as dialog boxes, buttons, and menus.

Launching a Program

IT Sentinel graphical programs are configured to run with default options that meet most user's needs. You can change certain aspects of a program's behavior by setting its preferences. See the section on the specific program in the Program Descriptions chapter of the *Sentinel Reference Guide* for details.

Note—The following instructions assume that the program is installed on your host system. If it is not, consult your system administrator or see the installation instructions provided with the software.

To open a program in the Windows environment, do the following procedure.

Procedure 4-1 Opening a Program under Windows

1. Log into a workstation that has an IT Sentinel installation and an available license.
2. Select Start > Programs > OPNET IT Sentinel <release>.

End of Procedure 4-1

To open a program in the UNIX environment, do the following procedure.

Procedure 4-2 Opening a program under UNIX

1. Log into a workstation that has an IT Sentinel installation and an available license.
2. At the operating system prompt, enter the program name, such as **itsentinel** or **spsentinel**. For example, to run IT Sentinel, type:

   itsentinel <*Return*>

End of Procedure 4-2

Exiting a Program

To exit an IT Sentinel graphical program, choose Exit from the File menu. You can also use the exit program capability of your operating system.

Editor Windows

The IT Sentinel program has several editors, each of which opens in its own window. You can open as many different editors as you like. The system window, which opens when you first launch IT Sentinel, is also an editor window.

Because editor windows are opened in the context of your workstation's window manager, they must be manipulated by the window manager's rules. See your system documentation for instructions on how to move or resize a program window.

The editor window includes five types of components:

- The *title bar*, which shows at least the name of the editor. The Project Editor title bar shows the name of the project, scenario, and subnet that is displayed.
- *The menu bar*, which contains pull-down menus giving access to all editor operations.
- *The button bar,* which contains buttons used to invoke certain operations.
- *A message display area*, which conveys information about editor operation.
- *The workspace*, where you do all work in the editor (create and edit models, run analyses, and so on).

The following subsections discuss these components.

This manual describes the standard IT Sentinel editor environment. However, the external tool support built into IT Sentinel allows value-added resellers (VARs) to make changes to this environment. If you are using a VAR-modified version of IT Sentinel, there may be additional menus or changes to the function of some operations. See the documentation provided by the VAR for information about such changes or additions.

Figure 4-1 IT Sentinel Editor Environment

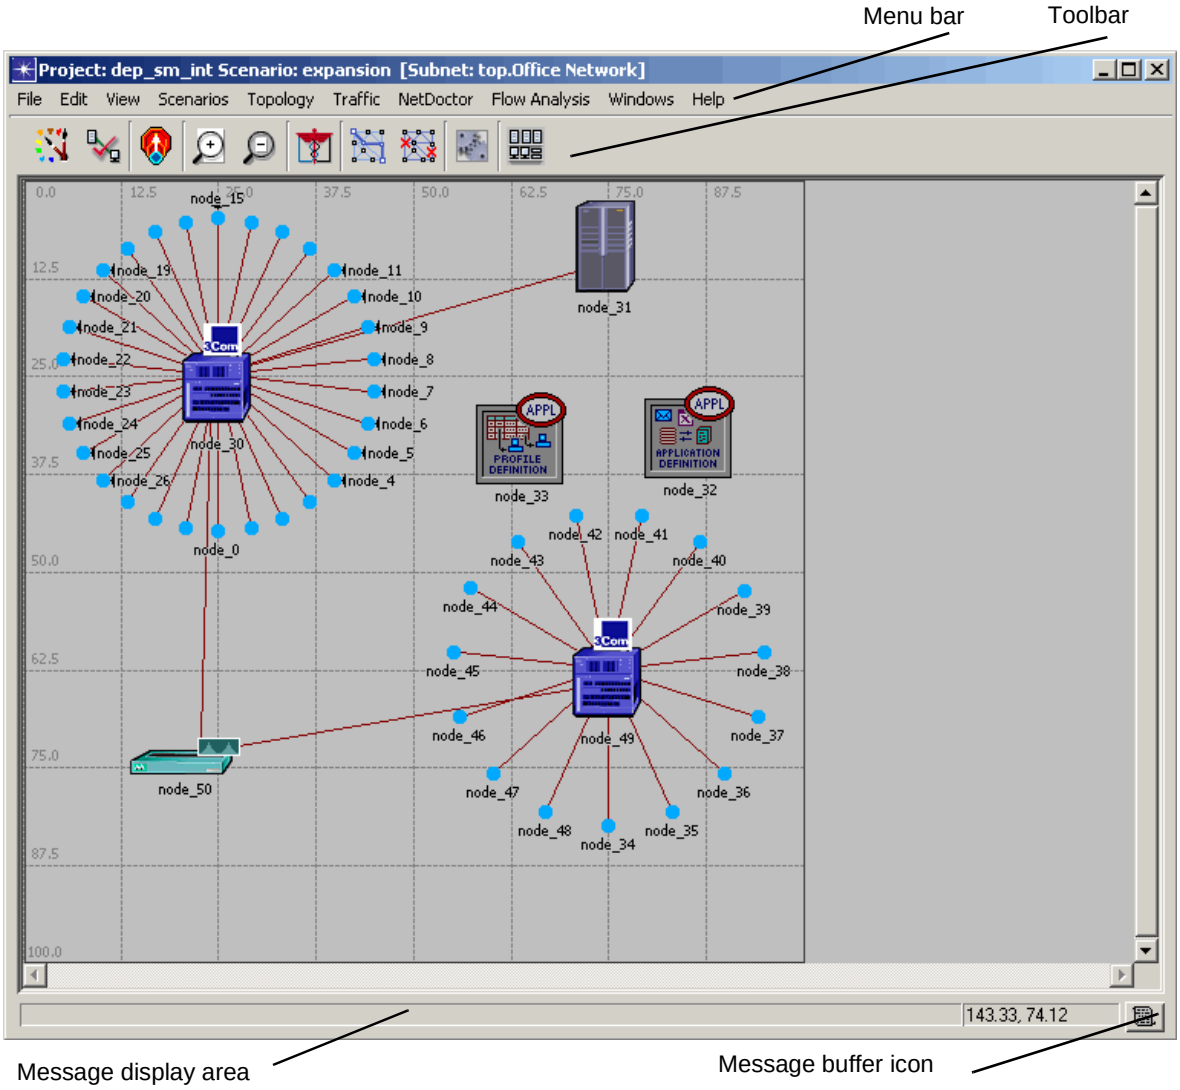

Menus

The menu bar provides access to IT Sentinel operations through conventional pull-down menus. Each menu contains *menu items* that invoke an operation, present information, or change the program display. The menus and available operations change, depending on the currently open editor, if any.

Shortcuts

Some of the menu items list a shortcut you can use to invoke the corresponding operation from the keyboard. For example, the Close item on the Project Editor's File menu has the shortcut "Ctrl+w", meaning you can close the window by typing <Control>+w. You can change the default shortcuts and define new ones as described in Editing Keyboard Shortcuts on page ISR-2-27 in the *Reference Guide*.

Pop-Up Menus

In addition to the menu bar menus, IT Sentinel has *pop-up menus* that appear when you right-click in various places within an editor window. These menus include many operations from the menu bar menus, along with some additional operations that apply to the context within which the menu was invoked.
IT Sentinel includes the following pop-up menus:

- Workspace pop-up menu—contains operations related to setting the workspace view, collecting results, and viewing results, as described under Workspace Pop-Up Menu on page ISU-6-35.
- Object pop-up menu—contains operations related to setting object properties, collecting results, and viewing results, as described under Object Pop-Up Menu on page ISU-6-36.

Tooltips

A context-sensitive tooltip appears when the cursor has pointed at an object for a few seconds. A tooltip explains the button or menu item and, for some operations, prompts for expected input.

Opening an Editor Window

Every IT Sentinel graphical program has one or more operations that open editor windows. These operations (generically called Open Editor operations) are invoked by various menu items, such as the New Project Editor item in the File menu. When invoked, an Open Editor operation opens an editor window and displays the menus for that editor.

Procedure 4-3 Opening a New Editor Window…

1. Choose New... from the File menu.

   ➥ The New dialog box displays.

2. Select the desired editor from the pull-down menu.

   ➥ The editor opens in its own window and becomes the active window.

End of Procedure 4-3

Procedure 4-4 Opening an Editor Window with a Specific Model

1. Choose Open... from the File menu.

   ➥ The Open dialog box displays.

2. Select the desired editor from the pull-down menu.

   ➥ The scrolling list changes to reflect the models that can be manipulated with the selected editor.

3. Select the desired model from the scrolling list.

   ➥ The desired editor and model are displayed in a new, active window.

End of Procedure 4-4

Controlling the Colors of the Interface

Two preferences, ui_colors.tool_back and ui_colors.tool_fore, control the color of window elements such as dialog boxes, backgrounds, and text.

Note—If an area's foreground and background colors match, the foreground elements will be invisible. Even closely-related colors can make foreground elements (such as black text against a dark-blue background) difficult to read.

Procedure 4-5 Setting an Element's Color

1. Choose Preferences from the Edit menu.
   ➥ The Environment Database window opens.
2. Click in the "Group" column header and choose Sort by Group.
   ➥ IT Sentinel sorts the preferences by group.
3. Scroll down the window until the UI-related preferences appear.
4. Click in the Value column of the preference you want to change.
   ➥ On UNIX platforms, the Color Chooser dialog box opens. See Color Chooser (UNIX only) on page ISU-4-21 for details on using the Color Chooser.
   ➥ On Windows platforms, the color chooser opens. See the Windows platform documentation for details on using the Windows color chooser.
5. Select the desired color from the color chooser and close the chooser.
6. In the Environment Database window, click OK.
7. Restart IT Sentinel to make the changed attribute take effect.

End of Procedure 4-5

Text Edit Pads

Text edit pads display view-only information about objects in an editor window, convey diagnostic messages, and allow you to enter information needed to complete an operation. In addition, they function as general-purpose text file editors and support text export for editing outside of the IT Sentinel program.

Functionally, text edit pads resemble the mouse-oriented text editors in most workstation window systems. They allow you to scroll through text; enter and delete characters; cut, copy, and paste selections; and read and write ASCII files. You can move and resize text edit pads using the same techniques as for windows.

Figure 4-2 Text Edit Pad

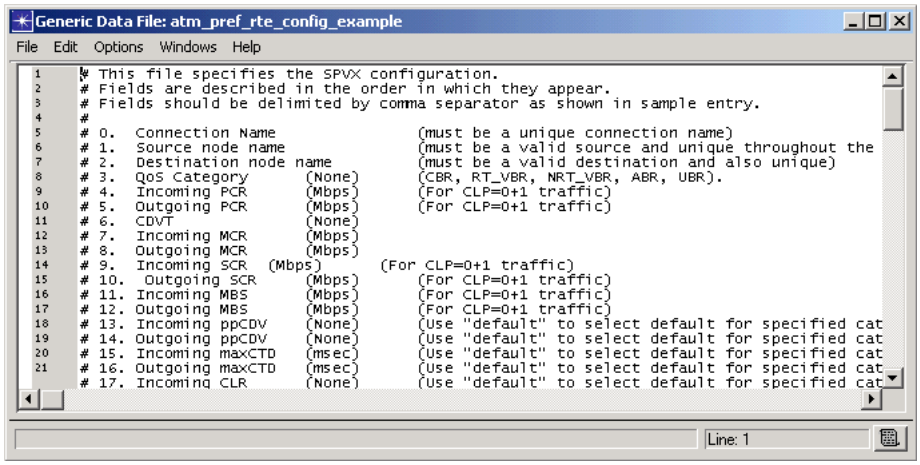

Text edit pads are large rectangular editing areas with scroll bars for accessing all regions of the pad. There is no limit on the number of lines. Options in the File and Edit menus allow various operations; buttons allow fast access to the most-commonly performed operations. Depending on the type of pad (view only or text entry), some operations may not be active. The line number is shown in the lower right corner of the text edit pad.

IT Sentinel allows use of external text editing programs (such as vi and emacs) as replacements for text edit pads. The preferences editor_keep and editor_prog support this capability. For more information on these attributes, see Program Descriptions in the *Reference Guide*.

The following sections summarize the operations available on each text edit pad menu.

File Menu (Edit Pad)

The File menu contains operations that relate to high-level functions such as opening, closing, and saving pads; importing or exporting contents; and printing. See also Table 4-7 File Menu Operations on page ISU-4-26.

Table 4-1 File Menu Summary

Menu item	Description	Reference
New...	Creates a new file in the current editor.	Opening a Text Edit Pad on page ISU-4-13
Open...	Opens an existing file in the current editor.	Opening a Text Edit Pad on page ISU-4-13
Close	Closes the current file.	Closing a Text Edit Pad on page ISU-4-14
Save	Saves the current file.	Saving the Contents of a Text Edit Pad on page ISU-4-15
Save As...	Saves the current file, allowing you to change the name.	Saving the Contents of a Text Edit Pad With a Different Name on page ISU-4-15
Page Setup...	Allows you to choose the page size and orientation for printing.	Printing a Text Edit Pad on page ISU-4-16
Print Preview	Shows an image of the text edit pad as it will be printed. From the print preview window you can choose the Page Setup or Print commands.	—
Print...	Allows you to specify the destination of a print command and the page range printed.	Printing a Text Edit Pad on page ISU-4-16
End of Table 4-1		

Edit Menu (Edit Pad)

The Edit menu contains operations that allow you to modify text and move quickly around the text edit pad.

Table 4-2 Edit Menu Summary

Menu item	Description	Reference
Undo Redo	Allows you to undo the last several operations or redo operations.	—
Cut	Copies the selected text to the clipboard and deletes it from the workspace.	Editing Text on page ISU-4-19
Copy	Copies the selected text to the clipboard.	Editing Text on page ISU-4-19
Paste	Replaces current selection with text from the clipboard.	Editing Text on page ISU-4-19
Select All	Selects all text in the text edit pad.	Selecting Text in a Text Edit Pad on page ISU-4-13
End of Table 4-2		

Options Menu (Edit Pad)

The Options menu contains operations that allow you to modify the appearance of text.

Table 4-3 Options Menu Summary

Menu item	Description	Reference
Show Line Numbers	Toggles the Line Numbers operation, which displays the number of each line in the left margin.	—
Drag And Drop Edit	Toggles the drag and drop operation, which allows you to move selected text by dragging it with the mouse.	Drag and Drop Editing on page ISU-4-21
Syntax Highlighting	Toggles the syntax highlighting operation. Syntax highlighting: • auto-indents lines • color-codes comments, literals, Kernel Procedures and other special types of text You can specify both the colors used and keywords that are color-coded with the code_coloring set of attributes. See the *Preferences* chapter of the *Reference Guide* for details.	—
Set Font	Changes the font that the editor uses to display all text in the file.	Setting the Font on page ISU-4-17
Save as Default	Makes the current Options settings (for line numbers display, drag and drop editing mode, syntax highlighting, and font) the default settings for new text edit pads.	—
End of Table 4-3		

Keyboard Commands

Text edit pads support many operations, most of which can be invoked by keyboard commands. For most of these commands, the text edit pad must be active for the command to be available.

Note—This tables describes the default keyboard shortcuts. You can change them by editing the operation files, as described in the System Environment chapter of the *Reference Guide*. The current definitions are shown next to the command on the menu.

Table 4-4 Text Edit Pad Keyboard Commands (Part 1 of 2)

Keystrokes	Operation
Movement Commands	
<arrow keys>	Moves the cursor one space in the indicated direction.
<Tab>	Inserts enough spaces to move the insertion point right to the next tab stop (located every two positions from the left margin).
Editing Commands	
<Backspace>	Deletes one character to the left of the insertion point, or delete all selected text.
<Control>+z	Undoes recent operation. Same as Undo button or menu option.
<Control>+y	Redoes operation. Same as Redo button or menu option.
<Control>+x	Cuts selected text to the clipboard. Same as Cut button or menu option.
<Control>+c	Copies selected text to the clipboard. Same as Copy button or menu option.
<Control>+v	Pastes the contents of the clipboard. Same as Paste button or menu option.
<Control>+b	Clears the selected text, but (unlike Cut) does not place the text in the clipboard. Same as Clear button or menu option.
<Control>+a	Selects all text in the text edit pad. Same as Select All menu option.
<Control>+f	Invokes the Find Text operation.
<Control>+g	Finds the next instance of the string in the Find dialog box.

Table 4-4 Text Edit Pad Keyboard Commands (Part 2 of 2)

Keystrokes	Operation
General Commands	
<Control>+n	Opens a new text edit pad. Same as New menu option.
<Control>+o	Opens the specified text edit pad. Same as Open menu option.
<Control>+w	Closes the text edit pad. Same as Close menu option.
<Control>+s	Save the contents of the text edit pad. Same as Save button or menu option. The text edit pad must be the active window. If another window is active instead, this command saves the contents of the active window.
<Control>+p	Print the contents of the text edit pad.
Navigation Commands	
<Control>+F2	Inserts or removes a bookmark at the current location. Same as Toggle Bookmark menu option.
F2	Jumps to the next bookmark in the file. Same as Next Bookmark menu option.
<Shift>+F2	Jumps to the previous bookmark in the file. Same as Previous Bookmark menu option.
<Control>+,	Moves the cursor to the specified line number. Same as Go to Line menu option.
End of Table 4-4	

Opening a Text Edit Pad

Most text edit pads open automatically in response to an operation being performed.

To open a new text edit pad at any time choose Open Edit Pad from the Edit menu of any editor or choose New from the text edit pad's File menu.

To open an existing file, choose Open from the text edit pad's File menu.

Selecting Text in a Text Edit Pad

You can use several methods to select text:

- Drag the cursor over the desired text.
- Double-click on a word to select that word.
- Triple-click to select the entire line.
- Click in the far left column of the text edit pad to select the entire line.

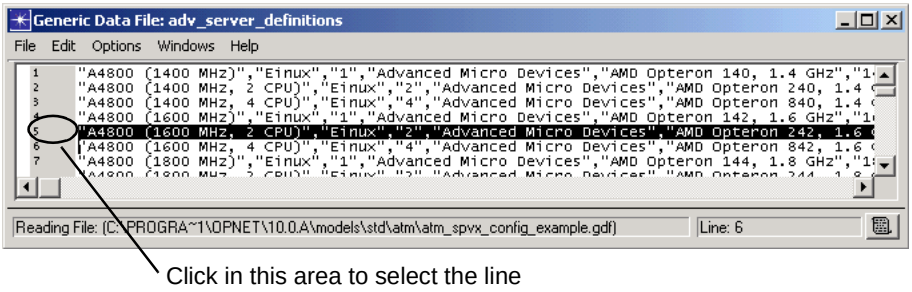

- Place the cursor on the first (or last) character of the text block you want to select. Hold down the shift key and place the cursor on the last (or first) character of the text block. The block demarcated by the two cursor clicks is selected.
- Choose Select All from the File menu to select the complete contents of the text edit pad.

Importing a Text File into a Text Edit Pad

You can import any text file into an open text edit pad.

Procedure 4-6 Importing a Text File

1 Choose Import from the File menu.

➥ An Open File dialog box opens.

2 Select the desired file, then click OK.

➥ The file is inserted into the text edit pad.

End of Procedure 4-6

Exporting a Text File to Disk

You can export a file from an open text edit pad to any directory.

Procedure 4-7 Exporting a Text File

1 Choose Export from the File menu.

➥ An Export dialog box opens.

2 Select the desired directory, type the file name into the text entry area of the dialog box, then click OK.

➥ The contents of the text edit pad are written to disk.

End of Procedure 4-7

Closing a Text Edit Pad

Closing a text edit pad does not save the contents. Use the Save operation before closing, if you want to save the contents.

Procedure 4-8 Closing a Text Edit Pad

1 Choose Close from the File menu.

➥ If you made changes, a confirm dialog box allows you to save them before closing. After you save (or decline to save) changes, the text edit pad closes.

End of Procedure 4-8

Saving the Contents of a Text Edit Pad

This operation saves the current text edit pad as a named file. Note that

- Some text edit pads are read-only and cannot be saved.
- Saving the contents of an Enter Execs or Exit Execs text edit pad does not save the contents to disk. This happens only after you save the process model.

Procedure 4-9 Saving a Text Edit Pad

1. Choose Save from the File menu.

   ➥ The Save dialog box opens.

2. Select the desired directory where the file will be stored.
3. Type the file name in the Save text entry area
4. Click on OK.

End of Procedure 4-9

Saving the Contents of a Text Edit Pad With a Different Name

The Save As... operation saves the current text edit pad with a different name.

Procedure 4-10 Saving the Current Text Pad Under a Different Name

1. Choose Save As... from the File menu.

   ➥ The Save As dialog box opens.

2. Select the directory where you want to store the file.
3. Type the file name in the Save As text entry area.
4. Click on OK.

End of Procedure 4-10

Printing a Text Edit Pad

You can print a text edit pad to

- The default printer specified in your operating system configuration.
- A different printer that you specify.

On a UNIX platform only, you can also send the contents of a text edit pad to:

- A file you name. The contents of the text edit pad will be converted to Postscript commands.
- A command. The contents of the text edit pad become the arguments of the command you specify.

Procedure 4-11 Printing a Text Edit Pad

1. Choose Page Setup… from the File menu.

   ➥ The Print Setup dialog box opens.

2. Specify the paper size and orientation, then click OK.

3. Choose Print… from the File menu or press <Control>+p.

   ➥ The Print dialog box opens.

4. Select the destination and page range to be printed, then click OK.

End of Procedure 4-11

Moving Within Text Edit Pads

The *insertion point* is a vertical bar cursor independent of the mouse cursor. It determines where text will be entered in a text edit pad. You can move the insertion point in several ways:

- Point at the desired location for the insertion point and left-click.
- Use the arrow keys to move the insertion point one space in the arrow direction.
- Use bookmarks to jump to a desired point in the pad. See Using Bookmarks on page ISU-4-17 for details.
- Choose the Go to Line operation from the Edit menu. Enter the desired line number and click OK. The line number of the current line is shown in the line number message display.

To view different parts of a text edit pad, use the scroll bars to change the view. You can also change the view by moving the insertion point. The text edit pad automatically scrolls to keep the insertion point (and surrounding text) visible.

Using Bookmarks

Set a bookmark on a specific line of text so you can move your cursor to it quickly. The bookmark appears as a rectangle in the left margin of the line. To place or delete a bookmark, put the cursor in the line-of-interest, then choose Toggle Bookmark from the Edit menu. After a bookmark is placed, use the Next Bookmark (F2) and Previous Bookmark (<Shift>+F2) menu options to move to the desired line of text.

Note—The bookmark cannot be copied and pasted. If you copy and paste a line that has a bookmark, use the Toggle Bookmark operation to set a new bookmark.

Setting the Font

This operation allows you to change the font used to display all text in the current file. You cannot use different fonts in the same file.

You can specify the type family (for example, Courier, Helvetica, Palatino), the display face (for example, bold, oblique), and size of the text display.

Procedure 4-12 Specifying the Display Font of a Text Edit Pad

1 Choose Set Font from the Options menu.

➥ The Font Chooser dialog box opens. A sample of the current font appears in the display area.

Figure 4-3 Font Chooser Dialog Box

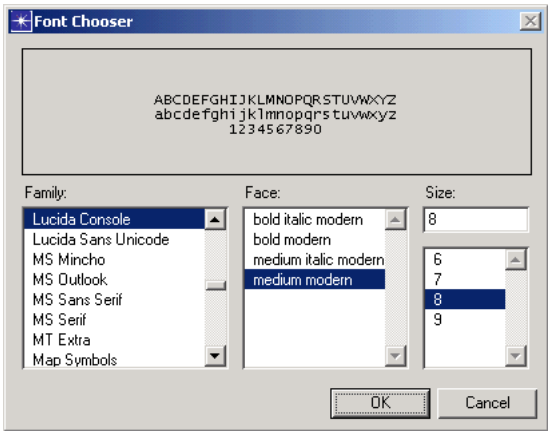

2 Select the desired type family, face, and size, then click OK. Some type families will have only certain faces or sizes.

➥ The text in the text edit pad changes to reflect the new font.

End of Procedure 4-12

Finding and Replacing Text

You can search for a text string. The search is conducted from the cursor location either forward or backward, as you specify. The cursor stops at the first occurrence of the specified string. You can configure your search to find complete words (“do” but not “undo”) and/or to make the search case-sensitive (“Project” but not “project”).

The Find Next operation on the Edit menu is equivalent to the Find Next button in the Find dialog box. Use this operation to repeat a defined search operation without reopening the Find dialog box.

Finding Text

To find text in an edit pad, do the following procedure.

Procedure 4-13 Finding Text

1 Choose Find… from the Edit menu.

➥ A Find dialog box appears.

2 Enter the desired text. Specify the search direction (forward or backward), the match criteria (whole words or case), and press Find Next.

➥ The cursor jumps to the first occurrence of the specified text, if it exists. If no matching string is found, the cursor does not move.

3 Click the Find Next button to repeat the find.

End of Procedure 4-13

Replacing Text

To replace text in an edit pad, do the following procedure.

Procedure 4-14 Replacing Text in a Text Edit Pad

1. Choose Replace... from the Edit menu.

   ➥ A Replace dialog box appears.

   Figure 4-4 Replace Dialog Box

   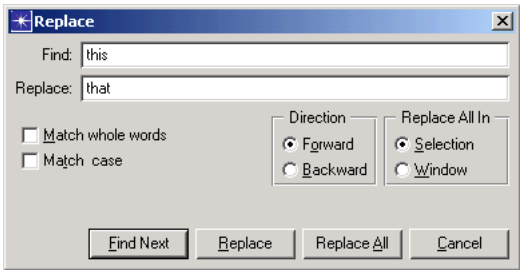

2. Enter the text to be found and the replacement text string. You can also specify the match criteria (whole words or case), the search direction (forward or backward), and the search region (selection or window).

3. Click on the button for the desired operation (find next, replace, replace all).

   ➥ The desired string is highlighted (if you specified find next) or replaced, if it exists. If no matching string is found, the cursor does not move.

End of Procedure 4-14

Editing Text

The cut, copy, and paste operations allow you to move or copy text within a text edit pad, from one pad to another, or between a text edit pad and another application. These operations store text in either a clipboard or a named file.

The clipboard used for editing operations stores the result of one cut or copy operation. Each time you cut or copy new text, the old text in the clipboard is replaced. All text edit pads share the same clipboard, so you can cut or copy text from one text edit pad and paste it into another. The contents of the clipboard are lost when you exit the program.

The clipboard is also used to exchange text between a text edit pad and other applications. In general, an application that can exchange text with **`xterm`** can also exchange text with an IT Sentinel text edit pad.

To cut or copy text to the clipboard, do the following procedure.

Procedure 4-15 Cutting or Copying Text to the Clipboard

1 Select the desired text.

➥ The selected text is highlighted.

2 Choose a cut or copy operation:

2.1 To cut the text, press <Control>+x, left-click on the Cut button or choose Cut from the Edit menu.

➥ The selected text is deleted from the text edit pad and placed on the clipboard.

2.2 To copy the text, press <Control>+c, left-click on the Copy button or choose Copy from the Edit menu.

➥ The selected text is placed on the clipboard. No change occurs to the text in the text edit pad.

End of Procedure 4-15

To paste text from the clipboard, do the following procedure.

Procedure 4-16 Pasting Text from the Clipboard

1 Place the insertion point where the text should be pasted, or select the text to be replaced.

2 Press <Control>+v, left-click on the Paste button or choose Paste from the Edit menu.

➥ The text in the clipboard is inserted or replaces the selected text. No change occurs to the text on the clipboard. This text can be inserted again and again.

End of Procedure 4-16

Drag and Drop Editing

With drag and drop editing, you can drag text from one location to another. This is quicker than using the cut and paste operation.

Procedure 4-17 Dragging and Dropping a Line of Text

1. Choose Drag And Drop Edit from the Options menu or verify that the option is already enabled.

2. Select the text you want to move. You can select a whole line by clicking in the far left column of the line.

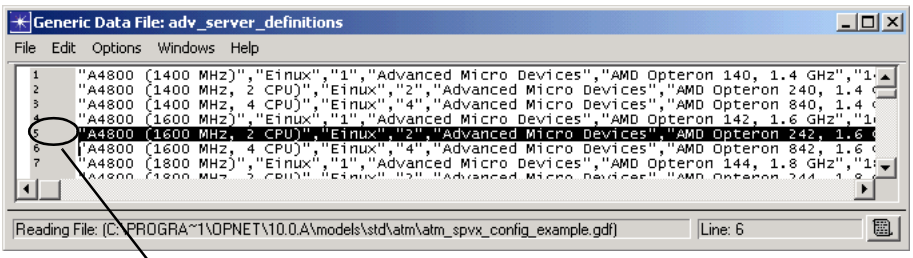

3. While pressing the left mouse button, drag the cursor to the new location.

   ➥ The text is moved to the new location.

End of Procedure 4-17

Color Chooser (UNIX only)

The Color Chooser appears if you are running the UNIX operating system. If you are on a Windows platform, the Windows color chooser appears instead. This section applies only to the Color Chooser. For help on the Window color chooser, see the platform documentation.

When you specify a color for a particular object, you select the color from the Color Chooser. The Color Chooser opens if you

- click on the Set Color button in a dialog box
- click on the color attribute of an object
- modify the value of a color-related preference, such as ui.colors.tool_back

You can select a default color or create and save a custom color. If you plan to apply the same color to different objects or attributes, save the color. You can then select the color from the Saved Colors palette and do not have to reselect it from the spectrum display each time.

Selecting a Color

To select a default or saved color, do the following procedure:

Procedure 4-18 Selecting a Default or Saved Color

1 Do an operation that requires a color selection.

➥ The Color Chooser opens.

Figure 4-5 Color Chooser Dialog Box

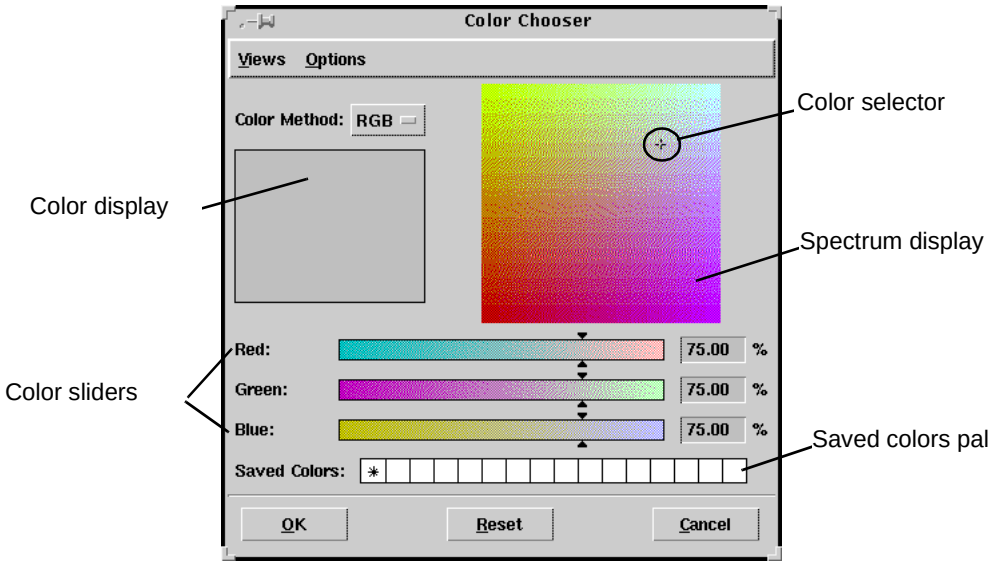

2 If you want the spectrum display to be a list of named colors instead of a color wheel or plane (a color plane is shown above), choose Named Colors from the Views menu. The Color Wheel/Plane option is the default.

➥ If you chose Named Colors, the spectrum display changes to a list of colors, in alphabetical order. You can rearrange the list by choosing Sort by Hue from the Options menu.

3 Select the desired color method (HLS, RGB, CMY, or Gray) from the pull-down menu.

➥ If you chose the Named Colors option, the spectrum display does not change.

➥ If you chose the Color Wheel/Plane option, the spectrum display may change to either a wheel (HLS or Gray color method) or plane (RGB or CMY color method).

4 If you are using the RGB or CMY color method, you may want to limit the spectrum display to the gamut of colors possible when only two (rather than all three) base colors are mixed. For example, if you are using the CMY color method, you may want to select either the Cyan-Magenta, Cyan-Yellow, or Magenta-Yellow gamut from the Options menu. When you select the Cyan-Magenta gamut, the spectrum display is limited to the colors possible when only those two base colors are mixed. To limit the spectrum display, follow these steps:

4.1 Choose Reform Slider Images from the Options menu.

4.2 Choose the desired display range.

5 Select the desired color:

5.1 To select a color from the spectrum display, left-click on it.

5.2 To select a saved color, double-click on it.

➥ The lower half of the color display changes to reflect your choice.

6 Click OK to close the Color Chooser.

End of Procedure 4-18

Creating and Saving a Color

To create and save a color, do the following procedure:

Procedure 4-19 Creating and Saving a Custom Color

1 Follow the instructions in Procedure 4-18 Selecting a Default or Saved Color on page ISU-4-22 to configure the spectrum display.

2 Select the color from the spectrum display that is closest to your desired color.

➥ The selected color displays in the lower half of the color display.

3 Adjust the color as desired by moving the indicators on the color sliders or by entering the desired percentage of the parameter. The parameters vary depending on the color method selected.

Table 4-5 Color Slider Parameters by Color Method

Color method	Adjustable parameters
HLS	hue, light, saturation
RGB	relative amounts of red, green, and blue
CMY	relative amounts of cyan, magenta, and yellow
Gray	relative amount of black and white
End of Table 4-5	

➥ The lower half of the color display reflects the adjustments you made.

4 If desired, adjust the dither method.

4.1 Choose Default Color Tolerances from the Options menu, and select the dither method desired. The table below explains the various color tolerances.

Table 4-6 Color Tolerances

Menu Selection	Dither Area (pixels)
Large Dither	8 x 8
Small Dither	2 x 2
Solid	No Dithering. On most displays, selecting too many colors that are not dithered limits the number of colors available for other uses.
Precise Large Dither	16 x 16
Precise Small Dither	4 x 4
Precise Solid	No Dithering. On most displays, selecting too many colors that are not dithered limits the number of colors available for other uses.
End of Table 4-6	

4.2 If you want to control dithering even more precisely, choose Color Tolerances from the Views menu, then specify a delta value that represents how close a color must match before dithering approximates the color.

5 Save the custom color.

5.1 Specify the cell that will store the custom color. Click on an empty (white) cell or click on an in-use cell to replace the current saved color with the new color.

➥ The asterisk appears in the cell that will store the custom color.

5.2 Choose Add Saved Color from the Options menu.

➥ The color in the lower half of the color display appears in the cell.

6 Click OK to close the Color Chooser.

End of Procedure 4-19

Reverting to the Previous Color

To revert to the previous color, do the following procedure:

Procedure 4-20 Reverting to the Previous Color

1 The color display shows the current color in the bottom half and the previous color in the upper half. To revert to the previous color, click in the upper half of the color display.

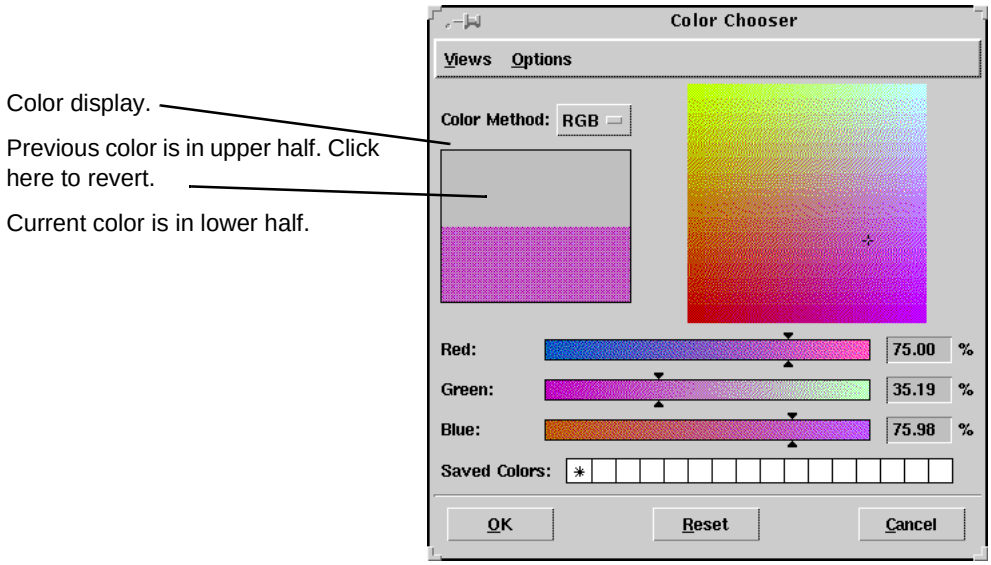

➥ The color display shows only one color, the previous color.

End of Procedure 4-20

System Menus

At startup—and anytime an editor is not open—IT Sentinel displays the System menus, which comprise the following menus:

- File
- Edit
- Automation
- License
- Windows
- Help

These menus contain items used to do system-level functions that are typically available in all editors. The following sections describe the operations found on these common menus.

Note—All menu operations do not appear at the same time. For example, on the File menu, New Project and Open Project appear only when no editor (such as the Project Editor) is active; at other times, New... and Open... appear on the File menu, and New Project appears in the Project Editor menu.

File Menu Operations

The File menu includes items that relate to high-level operations, such as opening and closing projects, saving scenarios, importing models, and printing graphics and reports. The following table summarizes the operations that can appear in this menu. The shaded operations appear only when an editor (such as the Project Editor) is active.

Table 4-7 File Menu Operations (Part 1 of 2)

Menu item	Description	Reference
New…	Creates a new file in the current editor.	—
Open…	Opens an existing file in the current editor.	—
Close	Closes the current project or file.	—
Save	Saves the current project or file. When saving a project, all scenarios within it are saved.	—
Save As..…	Saves the current project or file, allowing you to change the name.	—
Automation	Allows you to configure and run automation tasks	—
Delete Temporary Files		—

Table 4-7 File Menu Operations (Part 2 of 2)

Menu item	Description	Reference
Delete Projects	Select and remove a project and all related files from disk.	Delete Projects on page ISU-4-28
Model Files >		
Delete Model Files...	Removes individual project components, such as custom model lists, network models, etc.	—
Add Model Directory	Adds a new directory to the model directory (the mod_dirs preference).	Add to Model Directory on page ISU-4-28
Refresh Model Directories	Reloads the contents of all the model directories.	Refresh Model Directories on page ISU-4-28
Recent Files >		
<file type>	Displays a submenu listing the 5 most recently used files of the specified type. Select a listed file to open it. (There will be one of these menu items for each type of file used recently.)	—
Clear History	Removes all <file type> menu items and submenus.	—
Exit	Exit IT Sentinel	—
End of Table 4-7		

File Naming Conventions

IT Sentinel file names are not case-sensitive. You cannot change only the case of a file name; if you try to save a project that is already named “ABC.PRJ” as “abc.prj”, the file retains its original name.

You should not save two files with the same name, but different cases, in the same directory. In that case, only the file with the higher ASCII value will appear.

IT Sentinel file names can include any alphanumeric character and the underscore character (_). You should not include non-alphanumeric characters such as hyphens and periods in file names.

You should not use file names longer than 48 characters. If you use a longer file name, IT Sentinel truncates the name to 48 characters.

Delete Projects

The Delete Projects operation erases a project and its related files from disk storage. To delete a project, select it from the displayed list of available projects.

Note—There is no way to recover a project after you delete it.

Add to Model Directory

This operation adds a directory to your model directories list. After you select a directory, IT Sentinel adds the selected directory (and sub-directories, if specified) to your mod_dirs preference and refreshes its internal list of files to match the (updated) model directories.

Refresh Model Directories

The Refresh Model Directories operation updates a program’s internal list of files to match the IT Sentinel model directories. This may become necessary if you add or delete model files from the operating system command line while IT Sentinel is running, causing its knowledge of the model directories’ contents to become inaccurate.

When you choose Refresh Model Directories from the File menu, IT Sentinel searches the model directories and updates its internal list of accessible files. The operation takes a few seconds.

Closing a Program Window

Every GUI-based program provides an operation to close the program window and exit the program. To invoke this operation, choose Exit from the File menu.

Edit Menu Operations

The Edit menu includes items that enable you to manipulate text and objects, and to edit the preferences that control program operation. The following table summarizes the operations that can appear in this menu. The shaded operations appear only when an editor (such as the Project Editor) is active.

Table 4-8 Edit Menu Operations

Menu Item	Description	Reference
Undo	Undo the last several operations.	—
Redo	Redo the last undone operation.	—
Cut	Copies the selected object to the clipboard and deletes it from the workspace.	Editing Text on page ISU-4-19
Copy	Copies the selected object to the clipboard.	Editing Text on page ISU-4-19
Paste	Replaces current selection with the object from the clipboard.	Editing Text on page ISU-4-19
Delete	Removes object permanently and do not save to clipboard.	—
Select All in Subnet	Selects all nodes, links, and devices in the subnet.	—
Clear Model	Removes all objects in the workspace	—
Find Node/Link...	Finds, selects, and scrolls the editor window to display a node, subnet or link based on a full or partial name.	Find Node/Link Operation on page ISU-7-21
Select Objects...	Automatically selects or deselects network objects based on type, location in the network model, and attribute values.	Logical Object Selection on page ISU-9-15
Open Edit Pad	Opens an edit pad.	Opening a Text Edit Pad on page ISU-4-13
Preferences	Allows you to edit preferences for the program.	Preferences on page ISU-4-32
End of Table 4-8		

The Clipboard

The copy, cut, and paste operations use a *clipboard*, or hidden buffer, to temporarily store selected objects. The clipboard belongs to the editor environment rather than to a particular editor window. Thus, you can use the clipboard to copy objects within an editor window or from one editor window to another of the same type. Objects in the clipboard are lost when the program is exited.

Undo

Undo reverses selected operations: adding objects or moving, cutting, or pasting objects. You can undo an operation after you have scrolled or zoomed, but not after you have changed the attributes of the object that was operated upon.

If you have pasted a node with attached links and reattached the links, undo reverses the entire operation: the pasted node and the newly-reattached links are all removed.

Cut Selected Objects

The Cut operation deletes objects from an editor window and places them on the clipboard. Any dependent objects (such as links) associated with the cut objects are also deleted; any open dialog boxes associated with the cut objects are closed. You can also use this operation to delete items from the workspace.

You can undo one cut operation. When you undo a cut operation, copies of the cut objects remain in the clipboard.

Procedure 4-21 Cutting Objects

1. Select one or more objects to be cut.
2. Choose Cut from the Edit menu or type <Control>-x.

   ➥ The objects are copied to the clipboard and deleted from the editor window.

End of Procedure 4-21

Copy Selected Objects

The Copy operation copies objects in an editor window and places them on the clipboard. The original objects are not affected.

Procedure 4-22 Copying Objects

1. Select one or more objects to be copied.
2. Choose Copy from the Edit menu or type <Control>-c.

   ➥ The objects are copied to the clipboard. If the object copied is a node with connected links, the links are also copied. See Paste Clipboard Objects on page ISU-4-31 for a description of how the links are pasted.

End of Procedure 4-22

Paste Clipboard Objects

The Paste operation puts a copy of objects on the clipboard at a specified location within an editor window. The objects on the clipboard are unaffected and can be pasted repeatedly at different locations or in different editor windows. If the object you copied had connected links, the links are copied with the object. You can reconnect the links or not, as you want.

You can undo one paste operation.

Procedure 4-23 Pasting Objects

1 Choose Paste from the Edit menu or type <Control>-v.

➥ The cursor changes to four arrows and outlines of each object on the clipboard appear in the editor window.

2 Move the mouse until the object outlines reach the desired location.

➥ The cursor and outlines follow the mouse's movement.

3 Click the left mouse button.

➥ The objects are pasted in the specified location and the cursor resumes its usual appearance. If the name of an object being pasted is identical to that of an object in the editor window, a suffix is automatically added or incremented.

4 If the object being pasted had links and you don't want to reconnect them, right-click to end the paste operation. The object remains selected until you click elsewhere in the workspace.

5 If the object being pasted had links and you do want to reconnect them, follow these steps:

5.1 Move the cursor to each destination node and click. If you are pasting a subnet or connecting to a subnet, you may need to specify the source and destination nodes to be connected within the subnet.

5.2 Right-click at any time to end the paste operation. The object remains selected until you click elsewhere in the workspace.

End of Procedure 4-23

Clear Model

The Clear Model operation removes the currently displayed model from the editor window. It does not delete the model from disk.

Procedure 4-24 Clearing a Model from the Editor Window

1 Choose Clear Model from the Edit menu.

➥ If there have been no changes to the currently active model since it was last saved, the model is erased from the editor window. If there have been changes, a Confirm dialog box appears.

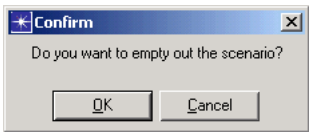

1.1 Click OK in the dialog box or press <Return>.

➥ The model is erased from the editor window.

End of Procedure 4-24

Preferences

Use this menu item to display and edit preferences, which control program operation, as described in the Preferences chapter of the *Reference Guide*.

Procedure 4-25 Editing a Preference

1 Choose Preferences from the Edit menu.

➥ A dialog box appears listing all the preferences you can edit. The title of the dialog box is the path to your environment database file.

Figure 4-6 Preferences Dialog Box

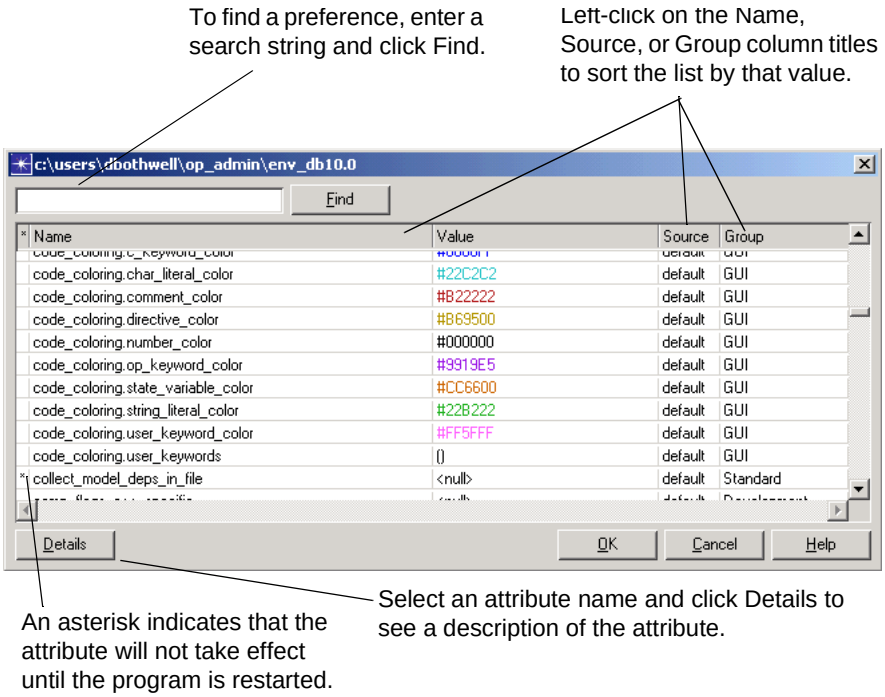

2 Type the first few letters in the name of the attribute you want to edit. If you wanted to edit the mod_dirs preference, for example, you would type `mo`.

➥ The list scrolls to the first attribute beginning with the typed letters.

3 Left-click in the Value column of the attribute to be edited and edit it using one of the following methods.

- For attributes with defined values (such as True and False), select the desired value.
- For attributes with open values, choose Edit… and enter the desired value.
- For attributes whose value is a list, edit the list in the dialog box that appears, as shown in the following diagram.

Figure 4-7 Dialog Box for Editing List Attributes

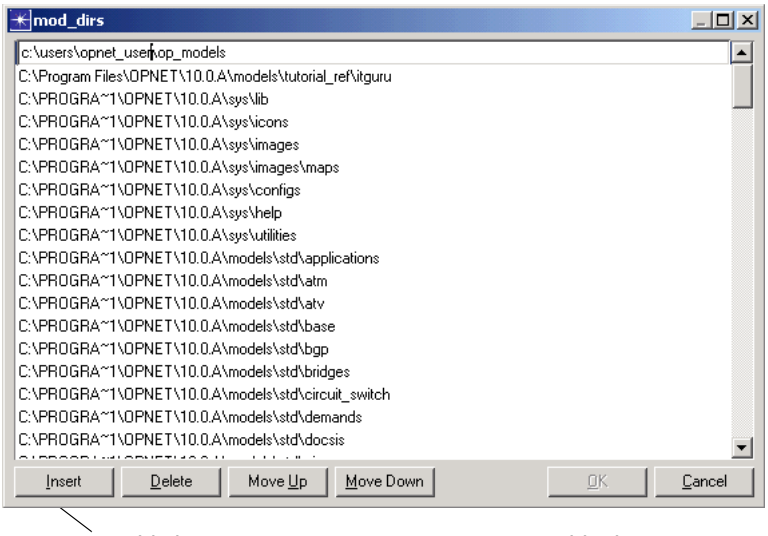

Use this button to copy a row or create a new blank row.

- To copy, select the row, making sure the text is highlighted, then left-click Insert.
- To create a new row, select a row, press <Return> to unselect any text, then left-click Insert.

Note—Some preferences (as identified in their Details listing) appear in the dialog box, but cannot be edited. To change the value of these command-line-only attributes, enter them as command-line flags when you invoke itsentinel.

---

4 Click OK to close the Preferences dialog box.

➥ The attribute value is changed in your environment database file. Most changes take effect immediately. However, some changes can only take effect when the program starts; if you have changed such an attribute, a dialog box will appear informing you of this.

Note—When the environment database file is first created, attributes are grouped by function and commented. After being edited, the file is rewritten as a simple alphabetical list of attributes. Comments are deleted.

End of Procedure 4-25

Automation Menu

The Automation menu contains operations that enable you to configure and run Sentinel automation tasks. The following table summarizes this operation.

Table 4-9 Automation Menu Summary

Menu Item	Description	Reference
Configure/Run Automation Tasks		Chapter 1 Overview on page AUG-1-1 of the Automation *User Guide*
Open Automation Log Manager		
Web—Open Report Server Home		
End of Table 4-9		

License Menu

The License menu contains an operation that allows you to open the License Manager. The following table summarizes this operation.

Table 4-10 License Menu Summary

Menu Item	Description	Reference
License Management	Opens the License Manager	License Manager chapter of the *Administrator Guide*
Product Modules	Opens Product Modules options	
End of Table 4-10		

Windows Menu

The Windows menu options allow you to hide or show specific windows. When you hide a window, the file is not closed or saved.

Table 4-11 Windows Menu Operations (Part 1 of 2)

Menu Item	Description	Reference
Previous Editor	Makes the previous editor the active window.	—
Circulate Editors	Moves among open editors.	—
Hide This Editor	Hides, but does not close or save, the current editor.	—

Table 4-11 Windows Menu Operations (Part 2 of 2)

Menu Item	Description	Reference
Hide Other Editors	Hides, but does not close or save, other editors.	—
Show All Editors	Shows all previously hidden editors.	—
System Window	Makes the system window the active window.	—
End of Table 4-11		

Help Menu Operations

The Help menu appears on the menu bar at all times. It provides access to context-sensitive help, the documentation, tutorials and other functions.

Table 4-12 Help Menu Operations (Part 1 of 2)

Menu Item	Description	Reference
Product Documentation	Opens the documentation.	Product Documentation on page ISU-4-37
Release Notes	Opens the release notes.	—
Tutorial	Opens the tutorial.	Tutorials on page ISU-4-37
Web – OPNET Home	Opens your browser to the OPNET home page.	Web – OPNET Home on page ISU-4-38
Web – Support Center	Opens your browser to the OPNET Support Center	Web – Support Center on page ISU-4-38
Web – Contributed Models	Opens your browser to the Contributed Models Depot	Web – Contributed Models on page ISU-4-38
Error Log >		
Open	Opens an edit pad with the contents of the error log.	Error Log on page ISU-4-38
Clear	Deletes the contents of the entire error log.	Error Log on page ISU-4-38

Table 4-12 Help Menu Operations (Part 2 of 2)

Menu Item	Description	Reference
Session Log >		
Open	Allows you to view the contents of the session log.	Session Log on page ISU-4-39
Clear	Deletes the contents of the entire session log.	Session Log on page ISU-4-39
About This Application	Displays information about the program (including release number) and the workstation environment.	About This Application on page ISU-4-39
End of Table 4-12		

Product Documentation

This operation activates IT Sentinel product documentation. The documentation contains all text and diagrams from the printed manuals, plus additional information about the software. The documentation uses a menu and hypertext system permitting easy navigation within the documentation. See Chapter 3 Using the Documentation Effectively on page DG-3-1 of the *Documentation Guide* for tips on navigating and finding information in the documentation.

To view the documentation, choose Product Documentation from the Help menu.

What's New in <Release Number>

This operation opens the What's New... menu. You can select to view the release notes for the current version IT Sentinel, the release notes for the previous release, or a web link that launches the browser and opens the OPNET Support Center menu page so you can select and view release notes from past releases.

To view the Release Notes menu, choose What's New in <Release Number> from the Help menu.

Tutorials

IT Sentinel comes with a tutorials to introduce you to basic program operation, as well as provide an overview of different protocols and advanced topics. The tutorial is provided in a small Acrobat Reader window that you can keep open on your workstation screen next to the IT Sentinel window, so you can do each lesson's tasks as you read about them.

The tutorial lessons are subdivided into topics with separate entry points, so that you can leave the tutorial at the end of any topic and easily resume.

To open the tutorial, choose Tutorials from the Help menu.

Web – OPNET Home

This operation launches your web browser to the home page of the OPNET web site.

Web – Support Center

This operation launches your web browser to the Support Center of the OPNET web site.

Web – Contributed Models

This operation launches your web browser to the Contributed Models Depot of the OPNET web site.

Error Log

This menu option contains two sub-options: Open and Clear. Open allows you to view your error log file (<*opnet_user_home*>/op_admin/err_log) in a text editor. When IT Sentinel encounters an error, it usually generates a message in this file.

This operation is similar to the program op_vuerr, which displays error messages via a command-line interface. See the description for op_vuerr (in the Program Descriptions chapter of the *Reference Manual*) for more information on error messages and error log files.

Note—Your error log file can become quite large with heavy use. You will probably want to clean out your error log periodically—especially if the log file takes several seconds or more to load into the text editor.

The Clear operation removes all existing messages from your error log file. You will want to do this periodically to remove obsolete messages and keep the error log manageable.

Clearing the log can also be useful when investigating and reproducing an error. You can clear the log and then reproduce the error; this ensures that all log entries relate to the error you are investigating.

Session Log

The Session Log contains all of the information generated during the current OPNET session. It is reset every time you restart OPNET. The Session Log menu option contains two sub-options: Open and Clear. The Open operation allows you to view your session log (<*opnet_user_home*>/op_admin/session_log) in a text editor.

The Clear operation removes all existing messages from your session log file. Unless your session is very long, you will probably not need to do this, as the session log is cleared each time you restart OPNET.

About This Application

The About This Application operation displays a window with information about the application and its execution environment. The following topics are covered:

- Environment information—provides information about the program's installation and environment, including current release and distribution level, host computer information, directory information, and your Group ID number, which is needed when contacting OPNET Technical Support.
- Notices—lists software licenses related to the OPNET product.

Procedure 4-26 Viewing Program Information

1 Choose About This Application from the Help menu.

➥ A view-only edit pad containing program information appears.

End of Procedure 4-26

Review the information in the edit pad, and close it when finished. See Text Edit Pads on page ISU-4-7 for instructions on using text edit pads.

5 Network Overview

A real-world network consists of physical sites such as workstations and switches (called *nodes*) that send, route, and receive information. IT Sentinel uses communication node objects to represent nodes. Real-world nodes communicate across *links,* which carry information over distances through electrical or fiber-optic cables. IT Sentinel represents these connections with communication link objects.

In complex communication systems, it is often necessary to view groups of nodes in a network as one entity. A group of related nodes and the links that connect them is called a *subnetwork*. In IT Sentinel, you can create a network hierarchy from subnetworks; this makes it easier to view and work with large networks.

This chapter describes how network objects work within an IT Sentinel network model. Procedures for creating objects and network models are in Chapter 6 Building Networks on page ISU-6-1 and Chapter 9 Customizing Models on page ISU-9-1. These chapters contain procedures that use operations available in the Project Editor.

Network Model Structure

A network model defines the overall scope of the system to be simulated, including its geographic scale, physical topology and logical configuration, its simulation and analysis parameter settings, and the statistics and report settings chosen for generating results. Each node in a network model serves to model a particular communications device or facility. There are no requirements or limits to the number of nodes represented in a network model.

The scope of a network can range from simple to complex. It might model one node, one subnetwork, or many interconnected nodes and subnetworks, because the structure and complexity of a network model typically follows those of the system being modeled. For example, a network with a star topology has a corresponding network model with one center hub node and several peripheral nodes connected to it with point-to-point links, as shown in the following figure.

Figure 5-1 Star Topology in Abstract and Network Model Representations

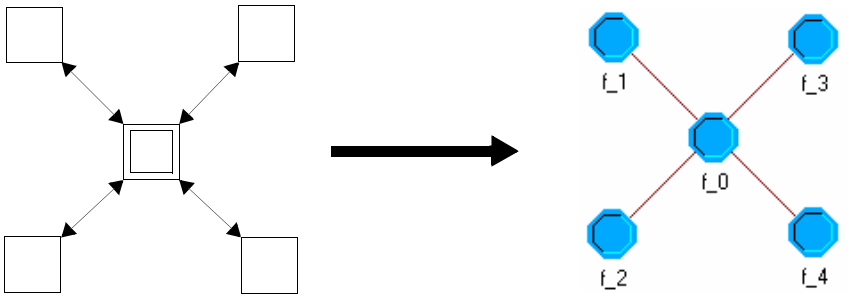

abstract

network model representation

IT Sentinel uses a project and scenario approach to modeling networks.

Projects

A *project* is a network design for which there might be several approaches. Each approach to implementing the network is a separate instance called a scenario.

Scenarios

A *scenario* is one instance of a network. Typically, a scenario presents a unique configuration of the network, with its own topology, protocols, applications, baseline traffic, and simulation settings. All projects contain at least one scenario.

Subnetworks

A subnetwork contains other network objects and abstracts them into one object. A subnetwork can encompass a set of nodes and links to represent a physical grouping of objects (for example, a local area network), or it can contain other subnetworks. Subnetworks within other subnetworks form the hierarchy of the network model. This hierarchy can be extended as required to model the structure of the network.

A subnetwork is considered the *parent* of the objects inside of it, and the objects are the *children* of the subnetwork. The highest level subnetwork in the network hierarchy is the *top subnetwork* or *global subnetwork*; it does not have a parent. Subnetworks can be created and interconnected within this top level or within other subnetworks.

Figure 5-2 Subnetworks Abstract the Network Components Contained Within Them

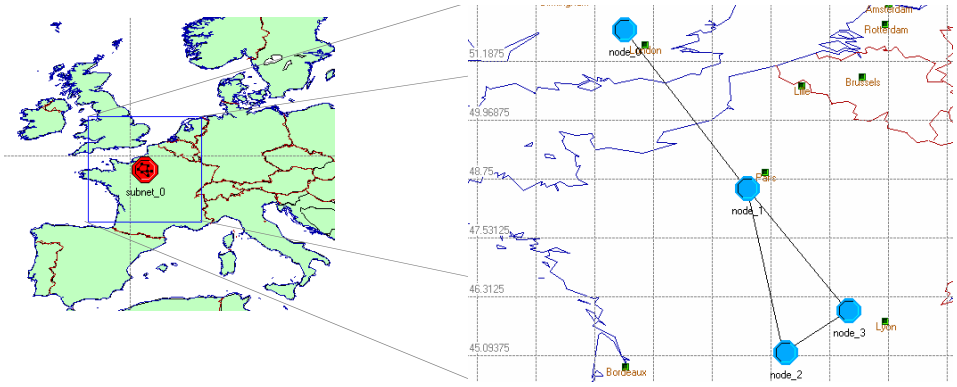

Subnetworks provide a powerful mechanism for manipulating complex networks by breaking down the system's complexity through abstraction. A large network with many components can be segmented into distinct parts based on the proximity, connectivity, or other architectural considerations of its elements. For example, subnetwork objects can represent campuses within a university. Within each campus, other subnetwork objects can represent each building or department, and so on.

Other than the objects it contains, the primary attributes of a subnetwork are its geographical position and size. However, these attributes can be ignored when a subnetwork is used strictly to abstract other network level objects or when subnetworks exist independently from each other.

Communication Nodes

Communication nodes exist within subnetworks and represent network devices that transmit and receive information. A node's behavior is determined by its node model. The node model specifies the node's internal structure (for example, its protocol stack) and can be specified through its "model" attribute. A node object is an instance of its node model. Distinct instances of the same node model operate independently of each other for purposes of analysis and simulation and can be configured with distinct configuration parameters, just as distinct pieces of equipment are independent of each other in a real network, although they may have identical capabilities. A network model can contain any number of communication nodes, using the same or different node models

You create node objects by selecting a node model from the object palette, as described in Creating Nodes on page ISU-6-21.

IT Sentinel supports two types of communication nodes: *fixed* and *LAN*.

Fixed Nodes

Fixed nodes model static network devices such as workstations, gateways, and servers. A fixed node is often connected to other nodes via one or more point-to-point or bus links. A fixed communication node cannot move during simulation.

Figure 5-3 Typical Fixed Nodes

LAN Nodes

LAN nodes are a special kind of fixed node that represent an entire Ethernet, FDDI, or Token Ring LAN (local area network), along with its aggregate traffic, as one entity. These nodes work in both switched and shared topologies, and can contain a variable number of workstations, as well as a server. LAN nodes

can be directly connected to any Ethernet, FDDI, or Token Ring node, except a hub, or any fixed subnetwork. LAN nodes can also be connected to objects with different data rates (for example, a 10BaseT LAN connected to a 100BaseT LAN) by using an intermediate router or switch.

Figure 5-4 Typical LAN Nodes

Switched LAN Node

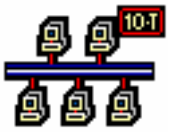

Shared LAN Node

Communication Links

Links enable nodes to communicate with each other. IT Sentinel supports two types of links: *point-to-point* and *bus*.

Each link type provides a fundamentally different type of connectivity. Point-to-point links connect one source node to one destination node; bus links connect each node of a fixed set to all other nodes in the set. Each type of link has attributes associated with it that control how the information is communicated between nodes. These can be grouped into the following categories:

- attributes that assign transmitter and receiver nodes
- attributes that determine how a packet is transmitted and when a packet is received
- attributes that control link behavior

A link's node attributes identify the nodes containing the transmitter and receiver to which the link is attached. IT Sentinel automatically assigns these values when a link is established between two nodes. You can change the automatic assignment by modifying the values of these attributes manually, as described in Setting Attribute Values on page ISU-9-4.

A link's "data rate" attribute specifies the rate at which information is transmitted over all communication channels in the link. This value should match the data rate set for each transmitter and receiver within the source and destination nodes. Because you might not always have access to transmitter and receiver data rates (unless allowed by the model developer), you can enter the special value unspecified for the link's "data rate" attribute. During simulation, this causes IT Sentinel to automatically assign the correct data rate for the link based on that set for the transmitter and receiver. If you specify this value when the "transmitter" and "receiver" attributes are also set to unspecified, IT Sentinel will set all three values to the link's default data rate.

You create link objects by selecting a link model from the object palette and designating the link's endpoints in the workspace. The Creating Links section describes this procedure.

Point-to-Point Links

A point-to-point link can be thought of as a bundle of one or more communication channels between the transmitter(s) and receiver(s) that it connects. Each communication channel acts independently of the others in the same link, as though it were defined in a separate and parallel point-to-point link. In an IT Sentinel network model, links are drawn between node and subnetwork objects. The connection between transmitters and receivers occurs automatically.

A *simplex point-to-point link* defines a connection from a transmitter in the source node to a receiver in the destination node. Packets are conveyed in that one direction. A *duplex point-to-point link* defines a pair of connections between two nodes, connecting a transmitter in each node to a receiver in the other. Packets can flow in both directions, from one node to the other. You can also use point-to-point links to connect networks, forming an internetwork. For example, information from transmitters in one network can be sent to receivers in another network. In this case, you would have two subnetwork objects connected via a link. Point-to-point links can also connect a subnetwork to one node.

A point-to-point link must be attached to nodes that contain point-to-point transmitters and receivers. Although it is possible in IT Sentinel to connect a node containing a point-to-point transmitter to a node containing a bus receiver, it is logically invalid and will result in simulation errors. To ensure that nodes and subnetworks are properly connected, the Project Editor has a *link consistency* operation. You can invoke this operation before executing a model to verify link connections and to avoid simulation errors. Instructions for using this operation are provided under Checking Link Consistency on page ISU-6-23.

Figure 5-5 Point-to-Point Links

Simplex Point-to-Point Link Duplex Point-to-Point Link

A packet completely arrives at a receiver after two delays have elapsed: the transmission delay and the propagation delay.

- The transmission delay is based on the transmitter channel's data rate and the size of the packet. For channel 0, the transmission delay matches the link's data rate. Each additional transmitter channel (if any) has its own data rate, which is specified by the model developer and may appear as a promoted attribute of the node model.
- The propagation delay is obtained from the link's "delay" attribute. This value applies to every communication channel within the link and can be specified as required to model the propagation delay across a link. The "delay" attribute can be specified as an absolute value or as *distance-based*, in which case the propagation delay is calculated from the distance between nodes and the link's propagation rate.

After a packet traverses a link, IT Sentinel computes the number of bit errors that occurred and determines whether they exceed the receiver's ability to correct them. The number of bit errors is computed from the value of the link's bit error rate (ber) attribute and the size of the packet. The number of bit errors that the receiver can correct is determined by the size of the packet and the receiver's error correction threshold (ecc threshold), which specifies the maximum number of errors per bit that can be corrected. For example, if a

packet's size is 1000 bits and the error correction threshold value is 0.002, then the receiver can correct the errors in the packet if the number of bit errors is 2 or less. If there are no errors or the receiver is able to correct those that occurred, the receiver accepts the packet. If there are too many bit errors, the receiver drops the packet and does not forward it.

6 Building Networks

IT Sentinel network models are based on three types of objects: subnetworks, nodes, and links. You work with these objects in the Project Editor. The Project Editor provides the resources you need to model the high-level components of a real-world network.

With the Project Editor, you can

- Create and edit network models
- Create derived models of nodes and links
- Customize the network environment
- Configure analyses for automated execution

The Project Editor contains a workspace for importing and reviewing network models that are being configured for automated analysis. Subnetworks and nodes are placed in the workspace as objects and depicted there as icons. Other objects, depicted as connecting lines, represent communication links between the nodes and subnetworks. Network objects are characterized by attributes that control how they behave within the overall model. The Project Editor provides operations for viewing these attributes (see Attributes on page ISU-9-1).

Building a Network Model

Building a network model consists of two steps:

1) Create a new project. When you create a new project, a new scenario is created automatically. You may also create additional scenarios for the project. See Projects and Scenarios on page ISU-6-2.

   You can create new projects in Sentinel two ways:

   a) Manually, when you define and test the automation tasks

   b) Automatically, as part of the operation of an automation task. Automation tasks can be configured to create new projects or use existing projects. For more information, see the Automation User Guide.

2) Import network topology and configuration. Sentinel can build the network from two primary data sources:

   a) Device Configuration Import—builds the network from device configuration files stored in external files. Typically, these files are exported from the devices or device management software.

   b) VNE Server Import—builds the network by connecting to VNE Server and obtaining topology and configuration. This also allows you to import link load data and end-to-end traffic flows along with device topology and configuration.

      See the Importing Network Data chapter for more information.

Projects and Scenarios

A project consists of one or more network scenarios. The scenarios may contain different versions of the same network or they may contain models of different networks. Each scenario corresponds to one network model. This section explains how to create, save, and close new projects. It also shows how to open existing projects.

To create a new project, do the following procedure.

Procedure 6-1 Creating a New Project

1 Choose New… from the File menu.

➥ The Startup Wizard appears to help you determine your first scenario's initial settings. Follow the steps indicated by the wizard.

End of Procedure 6-1

To save a project, do the following procedure.

Procedure 6-2 Saving a Project

1 Follow the appropriate step below:

- To save the active project, choose Save from the File menu. If the project has not been saved previously, you will be prompted to name the project.
- To save the active project under a new name, choose Save As... from the File menu.

➥ The project, including all its scenarios and data files, is saved.

Note—It is not possible to save an individual scenario. To save a scenario, save the project that contains it.

End of Procedure 6-2

To close a project, do the following procedure.

Procedure 6-3 Closing a New Project

1 Choose Close from the File menu.

➥ If the project has unsaved changes, a dialog box allows you to save those changes.

2 Choose Open… from the File menu.

➥ A list of available projects pops up.

3 Click on the name of the project you want to open.

➥ The project you selected appears in a new Project Editor.

End of Procedure 6-3

Using the Startup Wizard

The Startup Wizard helps you set the network environment, including

- topology
- scale
- map background
- grid spacing
- grid units
- model lists

If you select the “Use Startup Wizard” checkbox when creating a new project or scenario, the Startup Wizard will launch. Deselecting the checkbox causes the Startup Wizard to be ignored on future project or scenario creations. If you deselect the checkbox, and want to later use the Startup Wizard again, you must change the use_startup_wizard preference to TRUE. See Preferences on page ISR-4-22 in the *Reference Guide* for details about this attribute.

Working with Scenarios

A scenario has many different elements: always a network model, usually a probe model and a simulation configuration, possibly an analysis configuration file, a traffic import model, a Question/Answer database, a custom model list, an animation history, and a results file.

Scenarios are usually manipulated as a unit, but you may want to operate on individual elements. For example, you may want to pass a network to another user without the results or simulation configuration files.

New Scenario

To create a new scenario, do the following procedure.

Procedure 6-4 Creating a New Scenario

1. Display a new or existing project.
2. Select New Scenario from the Scenarios menu.

   ➥ The Enter Value dialog box displays with a default scenario name.
3. Enter a scenario name and left-click on OK, or just left-click OK to accept the default name.

   ➥ The Startup Wizard displays.

4 Enter values as prompted by the Startup Wizard to set the context of your scenario, or click on Quit to exit the Startup Wizard.

➥ IT Sentinel displays a workspace, with either the default context (if you did not specify a particular context in the Startup Wizard) or the context you specified.

End of Procedure 6-4

Duplicate Scenario

This operation places a duplicate of the selected scenario in the active project. It duplicates all elements of a scenario except simulation results.

To copy a scenario from one project to another, use the Import Components operation.

Procedure 6-5 Duplicating a Scenario

1 Display the scenario you want to duplicate.

2 Select Duplicate Scenario from the Scenarios menu.

➥ The Enter Value dialog box displays with a default scenario name.

3 Enter a scenario name and left-click on OK, or just left-click OK to accept the default name.

➥ The duplicated scenario displays.

End of Procedure 6-5

Manage Scenarios

The Manage Scenarios operation provides a concise overview of each scenario in your project and its status. Buttons and pull-down menus allow you to

- Rename scenarios
- Add new scenarios
- Duplicate scenarios
- Delete scenarios
- Re-order scenarios in the list

Changes are applied when you click OK to close the dialog box. No changes are applied if you click Cancel.

Procedure 6-6 Managing Scenarios

1 Select Manage Scenarios from the Scenarios menu.

➥ The Manage Scenarios dialog box for the current project displays

Figure 6-1 Manage Scenarios Dialog Box

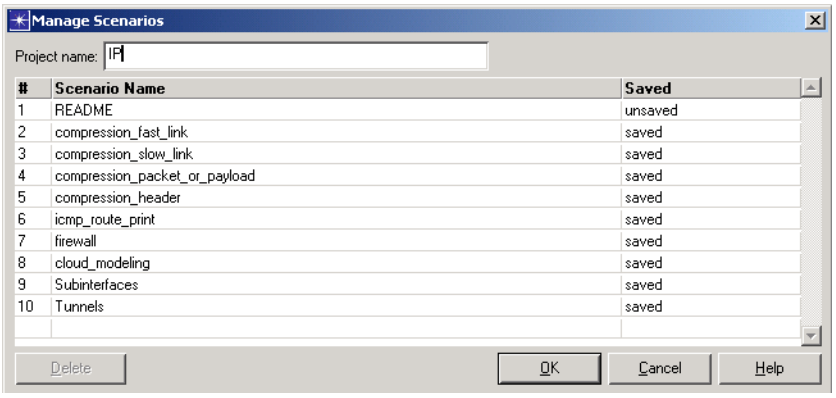

This dialog box uses the following conventions:

- Brackets mean a change to the scenario is requested and will be applied when you click OK.
- “Out of date” in a Results field means that the set of chosen results, simulation configuration, or network model has changed since the results file was created.
- “Uncollected” in a Results field means that a probe file exists, but a simulation has not been run.

2 Follow the instructions shown in Table 6-1 Manage Scenarios Operations on page ISU-6-7 for the desired operation.

3 Click OK to close the dialog box and store changes to disk.

End of Procedure 6-6

Table 6-1 lists the operations you can do from the Manage Scenarios dialog box.

Table 6-1 Manage Scenarios Operations

To…	Do this…
Rename a scenario	Select the scenario to be renamed and type the new name.
Add a new scenario	Left-click in the Scenario Name column of an empty row and select the <new> item from the menu.
Duplicate a scenario	Left-click in the Scenario Name column of the first empty row and select the duplicate <scenario_name> item from the menu.
Delete a scenario	Use either of two methods: • Select the scenario to be deleted and left-click on the Delete button. • Left-click in the Saved column and select <delete> from the menu. All files that comprised the scenario are deleted.
Re-order a scenario in the list	Left-click in the # column of the scenario to be re-ordered, then select the item that indicates the order desired. Select <leave> to keep the scenario in the same position.
End of Table 6-1	

Creating and Editing Network Objects

Object Palettes

You use object palettes to create objects in the Project Editor workspace, as described in sections Subnetworks on page ISU-6-18, Creating Nodes on page ISU-6-21, and Creating Links on page ISU-6-22. This section describes the object palette and how to open it.

Figure 6-2 Object Palette

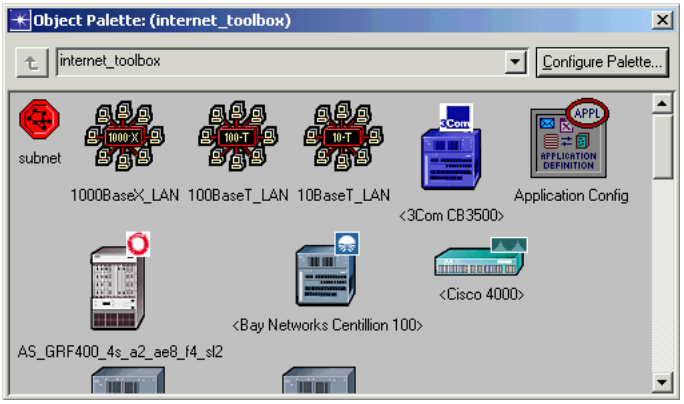

The object palette is divided into two areas. The top area contains buttons to configure and close the palette. The bottom area contains icons for a subnetwork and all available node and link models. By configuring the palette, you can limit the model icons displayed to those that match specified keywords (as discussed under Using Keywords on page ISU-6-10) or those included in a custom model list (as discussed under Using Custom Model Lists on page ISU-6-12).

The header bar of an object palette identifies the keywords or custom model list used by that palette. When multiple keywords are used, only the first keyword is listed, followed by an ellipsis (…) to indicate that there are more keywords. If the palette has not been configured, the header bar reads Keywords: <any>, indicating that all models are displayed.

Opening the Object Palette

To open the object palette, select Open Object Palette from the Edit menu. The object palette appears with icons for a subnetwork and all requested node and link models.

Note—If this is the first object palette you have opened during this session, there will be a short delay while your model directories are scanned and a list of available models is created.

Configuring Object Palettes

By default, an object palette displays all models in your model directories. If you have many models available, this results in a large palette that opens slowly and makes finding a specific model difficult. To make models easier to find, you can configure an object palette to display a subset of the available models.

For example, you may be creating an Ethernet-based network model. Normally, the Ethernet models would be mixed in with models supporting ATM, FDDI, and many other protocols. By configuring the palette you can cause it to display only the node and link models that support the Ethernet protocol.

There are two ways you can limit the models displayed in an object palette:

- Specify one or more keywords associated with the desired models
- Specify a custom list of the desired models

Both of these methods use the Configure Palette dialog box, which you open by left-clicking the Configure Palette… button. When you activate this button, the Configure Palette dialog box opens below the object palette and the text in the object palette's header bar changes to indicate that the palette is in Edit mode. In Edit mode you can specify keywords and create or specify custom model lists, but you cannot use the palette to create objects.

Figure 6-3 Object Palette with Configure Palette Dialog Box

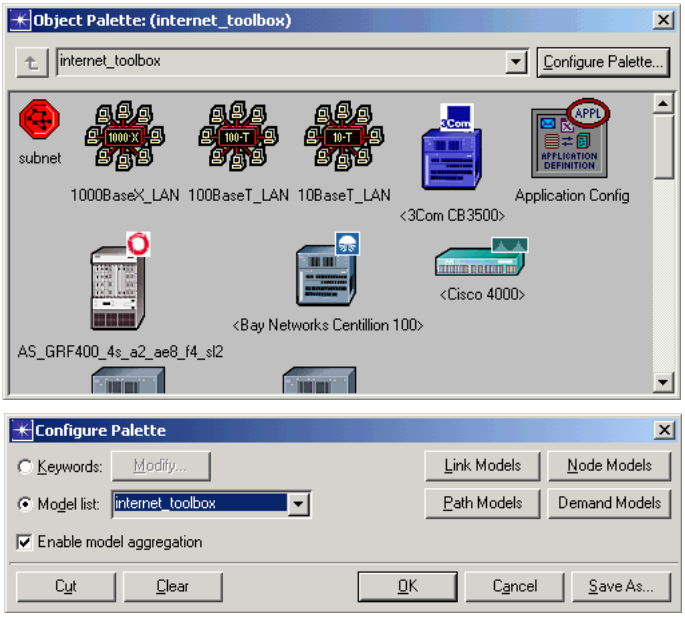

When you configure an object palette, the configuration is automatically applied to other object palettes you open later in the same

session. When you save a network model, the last configuration used within that model is saved with it and will be used the next time you open that network model.

To display sets of different models, you can open multiple object palettes at the same time, each with a different set of keywords or different custom model lists.

Using Keywords

All IT Sentinel-supplied node and link models have keywords associated with them. You can specify these keywords in the object palette to limit the number of models displayed. After you specify keywords, the object palette continues using them until you specify different keywords or a custom model list.

Note—The subnetwork object has no keyword and always appears in the object palette.

You can specify several keywords at once, in which case the object palette displays only models having all of the specified keywords. For example, the keywords “FDDI” and “link”, specified together, result in the display of FDDI link models only.

Some IT Sentinel models are assigned the reserved keyword “hidden” and are not displayed unless you explicitly specify this keyword. These models are hidden because they are very general in scope and often serve as parent models of IT Sentinel supplied derived models.

In addition to using the object palette, you can specify keywords from dialog boxes that include pop-up menus for model selection or via a preference when you run IT Sentinel. If specified at run time, the keywords will be used when you first open an object palette and thereafter until you change them. Run-time specification is described with the itsentinel “network_palette” preference in the Program Descriptions chapter of the *Reference Guide*.

Procedure 6-7 Specifying Keywords from the Object Palette

1. Left-click the Configure Palette… button.

   ➥ The Configure Palette dialog box opens.

2. Make sure the Keywords radio button is selected, then left-click the Modify… button.

   ➥ The Modify Keywords dialog box appears.

3. Left-click in the Status column to toggle the status of each keyword as needed:

   - Change the status of desired keywords to “set”.
   - Change the status of undesired keywords to “not set”.

4 When all keywords have the desired status, left-click OK.

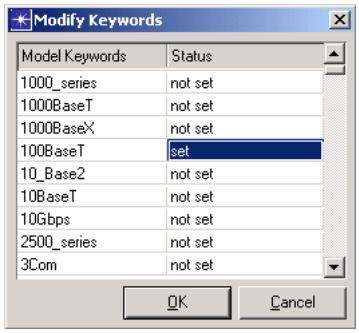

➥ The Modify Keywords dialog box closes and the object palette displays icons for only those models having all specified keywords.

End of Procedure 6-7

Using Custom Model Lists

A custom model list is a defined list of the models to be displayed in an object palette. It can contain any combination or number of models, regardless of keywords. You can create your own custom model list or use a pre-defined list. A custom model list can create a very small object palette that displays only the models you will use for a given network model.

To create a custom model list, you edit an object palette to display only those models you want and then save that set of models. The general process for creating a custom model list is as follows:

1) Open an object palette.

2) Open the Configure Palette dialog box, placing the object palette in Edit mode.

3) Use the Modify Keywords dialog box to select an initial set of models similar to those you want in the custom model list. [optional]

4) Use the Model List operations to delete models from or add them to the object palette.

5) Save the custom model list.

Creating a Custom Model List

To create a custom model list, do the following procedure.

Procedure 6-8 Creating a Custom Model List

1 Left-click the Object Palette toolbar button, or choose Topology > Open Object Palette.

2 Left-click the Configure Palette… button.

➥ The Configure Palette dialog box opens.

3 If desired, use Procedure 6-7 on page ISU-6-10 to specify one or more keywords for the object palette.

4 Left-click the Model List radio button.

➥ The buttons supporting custom model lists become active and the Modify… button becomes inactive.

5 Edit the object palette by adding or deleting models, as follows:

To add models (by name):

5.1 Left-click the Link Models or Node Models button.

➥ A Select Included Entries dialog box appears.

5.2 Left-click in the Status column to set each model's status to "included" or "not included", as desired, then left-click OK.

➥ Each model with the status "included" is added to the object palette.

To add models (by dragging):

5.1 Open an object palette that includes the model you want to add.

5.2 Click and drag the desired model's icon into the object palette being edited.

To delete models individually:

5.1 Left-click the models to be deleted.

➥ Selection markers appear on the models.

5.2 Left-click the Cut button.

➥ The selected models are removed from the object palette.

To delete all models, left-click the Clear button.

➥ All models are removed from the object palette.

6 Left-click the Save button and enter a name for the custom model list when prompted.

➥ The custom model list is saved to disk.

7 Left-click OK to close the Configure Palette dialog box.

End of Procedure 6-8

Specifying a Custom Model List

After you have created one or more custom model lists, you can specify which one to use in an object palette. As with keywords, you can specify a custom model list from dialog boxes that include pop-up menus for model selection or via an preference when you run IT Sentinel. If specified at run time, the custom model list will be used when you first open an object palette and thereafter until you change it. Run-time specification is described with the itsentinel network_palette preference in the Program Descriptions chapter of the *Reference Guide*.

To specify a custom model list, select the desired list from the model list pull-down menu. Icons of the selected list appear in the palette and the list name appears in the dialog box header.

Figure 6-4 Model List Pull-Down Menu

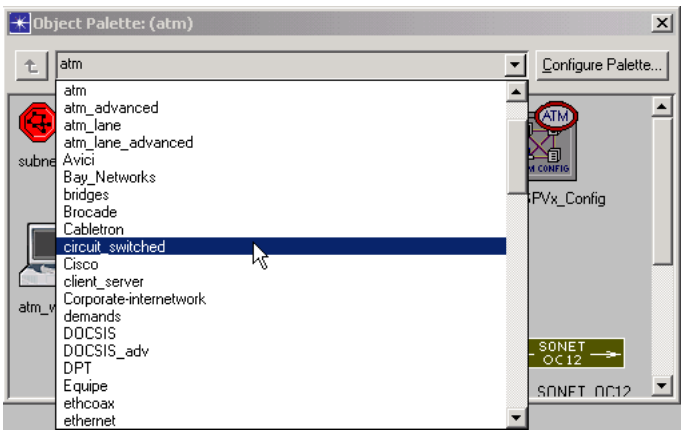

Grouping Models in the Object Palette

You can configure the object palette to display icons for groups of related models. This can make the palette less cluttered and easier to navigate, especially if the palette includes a large number of models.

Not all models in the OPNET model library are grouped. In general, vendor-specific node models (such as Cisco or Bay Networks) are organized into groups while most other models are not.

Figure 6-5 Grouping Models in the Object Palette

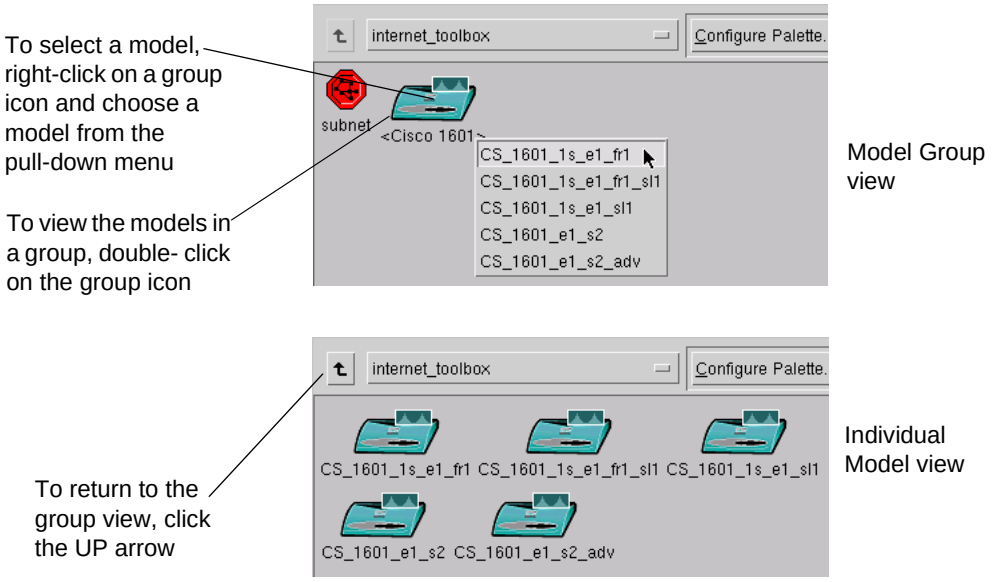

You can do the following in the object palette:

- To view models by group, check the Enable Model Aggregation check box in the Configure Palette dialog box. The palette aggregates all models that have an associated group. Angled brackets in an icon label (for example, <Cisco 1601>) indicate a model group.

  If Enable Model Aggregation is unchecked, the object palette displays all included models (one icon for each model) in one window.

- To select a model from a group, right-click on a group icon and choose the model from the pop-up menu. You can also click on a group icon and place it in the network; a pop-up menu opens and lists the models in that group.

- To view all models in a group, double-click on the group icon. The object palette "drills down" one level and displays all models in the group. To return to the previous level, click on the UP arrow in the upper-left corner of the palette.

A model's group is determined by the "System Object ID" attribute in the model's self-description. The object palette groups all included models with the same system object ID into one group.

Procedure 6-9 Viewing or Editing a Model's System ID

1 Open the model in the object palette.

2 Right-click on the model to open the Model Description dialog box. (Note: you must right-click on an individual model icon—not a model group icon—to open this dialog box.)

3 Click the Self-description button to open the self description for that model.

4 Right-click on the core icon to open the “core Self-Description” dialog box. The System Object ID is listed in the Characteristics table of this dialog box.

End of Procedure 6-9

Viewing Model Information

Every model has information associated with it that gives the model’s keywords, supported object types, a summary of its attributes, and a description of its use. The description is especially important, because it can tell you with what other types of objects the model can or cannot be used. This information can help you build valid network models and avoid simulation errors. It is a good practice to view a model’s description before using it.

You can view a model description by

- Selecting the View Node Description operation from the Node pop-up menu (or View Link Description from the Link pop-up menu). This displays the text description of the *model associated with the selected object*, in an edit pad that can be copied, printed, or written to a file.
- Opening the Model Description dialog box from within the object palette. This dialog box displays information on the *model itself*. Buttons in this dialog box allow you to display the parent model, derive a new model based on the existing one, display the properties of attributes, and do other useful functions.
- Reviewing the documentation in the *Reference Guide*.

Viewing a Node/Link Description

To view a node or link description, right-click on a node or link to view the Node or Link pop-up menu. Choose View Node/Link Description to view descriptive text in a text edit pad.

Figure 6-6 Model Description

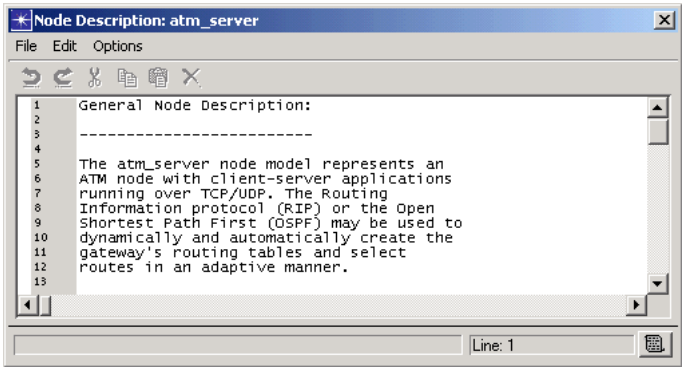

Viewing a Model Description

To view a model description, right-click on the icon of an individual model in the object palette.

Figure 6-7 Model Description Dialog Box

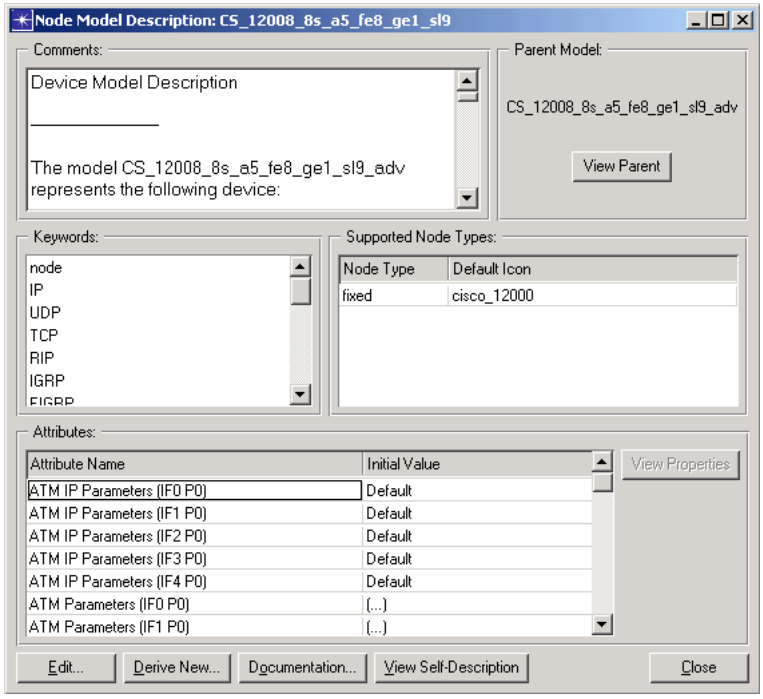

The Model Description dialog box contains read-only fields for the model information and five buttons for viewing additional information or doing other operations. The following table lists the fields and buttons.

Table 6-2 Model Description Dialog Box Contents

Label	Description
Attributes	Lists the name and initial value of the model attributes that will appear on an object created with this model.
Comments	Displays comments that describe the model and provide hints for using it.
Close	Closes the dialog box.
Keywords	Lists the keywords associated with the model.
Node (or Link) Documentation	Shows the Help Card file associated with the model, and the tool tip presented by the model.
Parent Model	Identifies the model's parent model. If the model is a base model, this area will read ***(none)***.
Self-description	Opens a self-description window for viewing model details.
Supported <*object*> Types	Lists the types of objects supported by the model and the icon used to represent each type when a new object is first placed in the workspace.
View Parent	Summons a Model Description dialog box for the model's parent model. This button is inactive if the model is a base model.
View Properties	Summons a view-only version of the Attribute Properties dialog box (as described in Attribute Properties Dialog Box on page ISU-9-28). You must select an attribute name in the Attributes table to make this button active.
End of Table 6-2	

Subnetworks

A subnetwork is typically the first object you will create when building a new network model. The following procedures describe how to create subnetwork object, view its contents, and return to its parent subnetwork.

Procedure 6-10 Creating a subnetwork

1. If necessary, open an object palette.
2. Select the subnet icon in the object palette.

➥ A square icon outline and movement cursor appear over the ***subnet*** icon.

3 Drag the icon outline to the desired location in the Project Editor workspace, then release the mouse button.

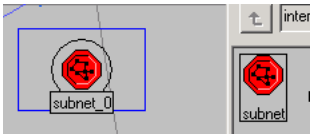

➥ IT Sentinel draws the subnetwork in the workspace and assigns a default name. The subnetwork icon is enclosed by a rectangle called the *subnetwork extent*. The extent defines the location and size of the subnetwork.

4 If desired, reposition the cursor and left-click to create additional subnetworks.

Note—The object palette might be set to create only one object, depending on its creation mode. If so, repeat steps 2 and 3 to create additional subnetworks (ignore steps 4 and 5). For details on creating single vs. multiple objects, see the creation_mode preference in Program Descriptions of the *Reference Guide*.

5 When you finish creating subnetworks, right-click to end the operation.

End of Procedure 6-10

Procedure 6-11 Entering a Subnet

1 Double-click on the desired subnetwork or right-click on the subnetwork and choose Enter Subnet from the pull-down menu.

➥ The subnetwork's contents, or *subnet view*, appear. You can place objects within this view or extend the hierarchy by creating additional subnetworks.

The only difference between double-clicking and the Enter Subnet operation is that double-clicking also deselects any other selected objects.

End of Procedure 6-11

Displaying the Parent of a Subnetwork

Subnetworks can be nested to create a hierarchy of network levels. After working within a lower-level subnet, you can return to the next higher level (the parent subnet) as follows.

Procedure 6-12 Returning to the Parent Subnetwork

1 From the View menu or the Workspace pop-up menu, choose Go To Parent Subnetwork.

2 If there are multiple levels of subnet hierarchy, repeat step 1 as needed.

End of Procedure 6-12

Setting Background Properties

The Set Background Properties operation (View > Set Background Properties) lets you redefine background properties of the current subnetwork. See Subnetwork Background Properties on page ISU-7-2 for more information.

Expanding a Subnet

Expanding a subnet displays the contents of the subnet at the parent subnet level. You see the area covered by the lower-level subnet in the context of the larger geography covered by the parent subnet. This is helpful when two subnets cover the same geographic area or cover areas that partially overlap, because you can clearly see the interaction among all objects. Say, for example, that you create two different subnetworks, placing a jammer in the first. Assume further that the jammer actually affects transmissions in the second subnetwork. This effect is not obvious when you look only at the second subnetwork but becomes clear when you see the contents of both subnetworks in the larger geographic context.

When you expand a subnet, you see the objects in it, the links leading to the subnet (reattached to the appropriate objects in the subnet), and the subnet extent, as determined by the subnet object's "x span" and "y span" attributes. While the subnet is expanded, you can

- Drag the subnet as a whole to a new location. The objects in the subnet move with it as you drag.
- Right-click on a object to display the Object pop-up menu, and do any of the operations in the pop-up menu.
- Attach links to an object in the subnet.

To move, cut, copy, or paste an object in the expanded subnet, you must actually be in the subnet. Double-click on the subnet's surrounding box (the subnet extent) to open the subnet in a new window.

Figure 6-8 Two Expanded Subnets

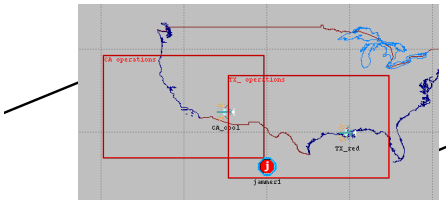

Collapsing a Subnet

This operation restores selected subnets to their typical collapsed state, in which they are represented by one icon. To select an expanded subnet, click on the subnet extent. See Expanding a Subnet on page ISU-6-20 for details on expanding subnets.

Figure 6-9 Expanded and Collapsed Subnets

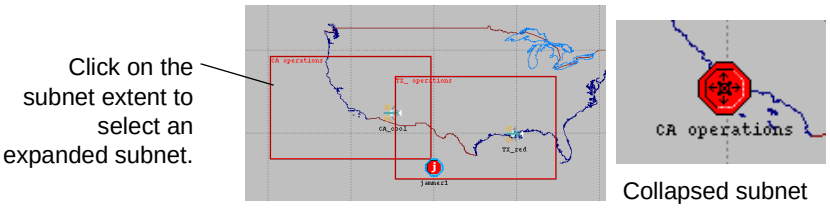

Creating Nodes

Unlike subnetworks, nodes (and links) must have an associated model to describe how they behave in the network. These models are represented by icons in the object palette. When you drag a node model's icon into the workspace (or select a link model's icon to connect nodes), you create an *instance* (copy) of that model. The node (or link) is depicted in the workspace as an object and its "model" attribute is automatically set to the one chosen from the palette. In addition, the object's attributes are updated to reflect the model specification, as described in Attributes on page ISU-9-1.

You can add node objects to the network model by dragging them from the object palette.

Procedure 6-13 Creating a Node using the Object Palette

1. If necessary, open an object palette and configure it to limit the display to only those models of interest.

2. Left-click on the model that you want to use for the object.

   ➥ A square icon outline and movement cursor appear over the icon.

3. Drag the icon outline to the desired location in the Project Editor workspace, then release the mouse switch.

   ➥ IT Sentinel creates the node in the workspace and gives it a default name. The node's "model" attribute is set to correspond to the model icon selected from the object palette.

   The default name consists of a text string (such as "node" for a fixed node or "l" for a link), followed by an index that is incremented each time an object is created. You can rename an object by editing its "name" attribute, as described in Setting Attribute Values on page ISU-9-4.

4 If desired, move the cursor to a new location in the workspace and left-click to create another instance of the node.

Note—The object palette might be set to create only one object, depending on its creation mode. If so, you must repeat steps 2 and 3 to create additional subnetworks, and can ignore steps 4 and 5. For details on creating single vs. multiple objects, see the itsentinel “creation_mode” preference in the Program Descriptions chapter of the *Reference Guide*.

5 When you are finished creating nodes, right-click to end the operation.

End of Procedure 6-13

Creating Links

Links are used to create communication channels between nodes. You create links by selecting a link model from the object palette and then graphically drawing the link in the workspace. You can create multiple-segment links and move them using the mouse (the end segments change length to reflect the new location).

Procedure 6-14 Creating a Link…

1 If necessary, open an object palette and configure it to limit the display to only those models of interest.

2 Left-click a link model icon in the object palette to select it (selection markers will not appear).

3 To create a point-to-point link:

3.1 Left-click the subnetwork or node that will transmit information.

3.2 Left-click on any intermediate points (if you want the link to follow a crooked path).

3.3 Left-click the subnetwork or node that will receive information.

➥ IT Sentinel draws the link between the two objects.

4 To create a bus link:

4.1 Left-click in the workspace where the bus link will begin.

4.2 If the bus will have any turns, left-click at each bend in sequence from the beginning of the bus.

4.3 Left-click the mouse while pointing at the location where the bus will end.

➥ IT Sentinel draws the bus link in the workspace. You will now have to create fixed nodes and taps to place around the bus.

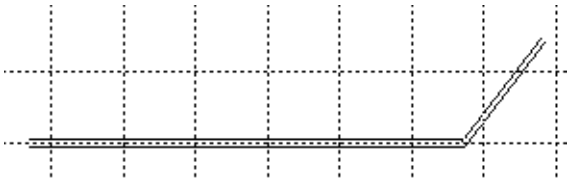

5 To create a bus tap:

5.1 Left-click the point on the bus where the tap will connect.

5.2 Left-click the node to be attached to the bus.

➥ IT Sentinel draws the bus tap in the workspace.

For all link types, when a subnet is selected as a link endpoint, IT Sentinel prompts you for the name of the node within the subnetwork which will be connected to the link.

End of Procedure 6-14

Checking Link Consistency

For a link between two objects to be consistent, the packet formats and data rates of the transmitter and receiver within the objects must match those defined for the link. When you connect two objects with a link, IT Sentinel automatically tries to create a consistent connection. However, this may not be possible if the models used by the objects are incompatible. For example, although it is graphically possible to connect two Ethernet nodes via an X.25 link, it is not logically correct and would result in an inconsistent link and subsequent simulation errors. The data rates defined for Ethernet and X.25 differ, and the packet formats may not be supported for the link.

The Verify Links operation lets you verify that links are properly connected before executing a simulation. This helps prevent simulation runtime errors by confirming that all links in the network model can transport packets.

When invoked, the Verify Links operation checks each point-to-point and bus tap link to confirm that

- the link connects to the proper number of transmitters and receivers in the attached nodes.

- the data rates of each transmitter and receiver match the data rate of the link. (However, links whose data rate is promoted or set to unspecified need not match.)
- the packet formats defined for each transmitter and receiver have at least one format in common with the link. The common format can be a specific value or unformatted.

The data rate is the value specified for an object's "data rate" attribute. The packet format is the value specified for an object's "packet formats" attribute. For more information about these attributes, see the Preferences chapter in the *Reference Manual*.

Note—Bus links are not checked for consistency because of the variable nature of the connections they support.

Procedure 6-15 Verifying Links

1 With a model in the editor window, choose Verify Links... from the Topology menu.

➥ The Check Links dialog box appears.

2 Select the Verify links radio button and left-click OK.

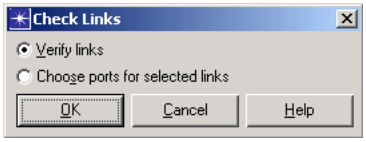

➥ IT Sentinel pauses for a moment while it checks link consistency. Each inconsistent link, if any, is selected and identified with an "X" mark both in the Project Editor.

3 Double-click the "X" mark.

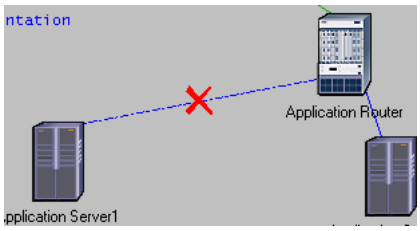

➥ A dialog box appears with information about the incorrect link.

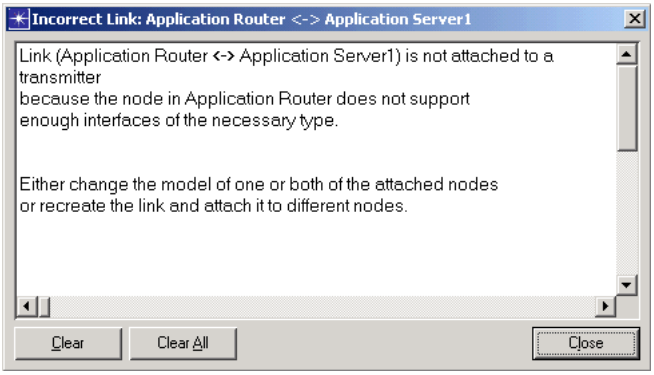

4 Based on the information in the dialog box, you may have to reassign a node's "model" attribute or use a different type of link to correct the problem. The following table lists common problems and solutions related to link consistency.

Table 6-3 Common Link Consistency Problems & Solutions

Problem	Solution
Inconsistent data rates between the transmitter, receiver, and link.	Compare the "data rate" attribute value of the link with that of each node. The values should match. Note that the "data rate" attribute may not be visible unless it was promoted from lower model levels. Instructions for viewing link and node attribute values appear in Attributes on page ISU-9-1.
Link inconsistent due to a missing bus transmitter.	Check the *Model Guide* or the Model Description dialog box to see what kind of transmitter and receiver are present in the objects connected to the link. Either change the type of link used to one that supports the transmitter and receiver present, or change the connected nodes to ones that contain a transmitter and receiver supported by the link.
Link inconsistent due to a missing bus receiver.	
Link inconsistent due to a missing point-to-point transmitter.	
Link inconsistent due to a missing point-to-point receiver.	
End of Table 6-3	

5 Click the Clear button to remove the inconsistency mark over the current link or click the Clear All button to remove all inconsistency marks.

You can choose to leave the marks visible. In this case, IT Sentinel will remove them automatically the next time you invoke the Verify Links operation, assuming the links are then consistent.

Note—If you change the "model" attribute of a node, it is possible that some of the links attached to that node may become inconsistent. It is always a good idea to check the consistency of network links whenever any network topology changes are made.

End of Procedure 6-15

Choosing Ports

When you create a link, IT Sentinel tries to connect it to *ports* in the source and destination nodes that will be consistent with the link's packet formats and data rate, as discussed under Checking Link Consistency on page ISU-6-23.

If IT Sentinel cannot find ports consistent with the link, the link consistency check will fail. In this case, you can take any of several actions, using the techniques described in Table 6-3 Common Link Consistency Problems & Solutions on page ISU-6-25:

- Remove the link and replace it with one that matches the nodes.
- Change the link's "model" attribute to one that supports the packet formats, data rate, and channel count required by the connected nodes.
- Change the values of the link's attributes to support the packet formats, data rate, and channel count required by the connected nodes.
- Change the node's "model" attribute to one that supports the packet formats, data rate, and channel count required by the inconsistent link.

After making one of these changes, you can manually cause IT Sentinel to reassign ports for consistency by activating the Choose Ports operation.

Procedure 6-16 Choosing Ports

1. Confirm that only links for which you want to reassign ports are selected. (Bad links are selected automatically by the Verify Links operation).
2. Choose Topology > Verify Links.

   ➥ The Check Links dialog box appears.
3. Select the Choose Ports for Selected Links radio button and left-click OK.

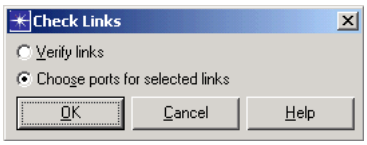

➥ IT Sentinel searches for consistent transmitters and receivers in the source and destination node or subnet of selected links and makes new assignments based on the results of the search. If a link is marked with an “X”, but is made consistent by the reassignment, IT Sentinel removes the “X”.

If there is still an inconsistent link, the “X” remains and you should try another of the previously mentioned solutions.

End of Procedure 6-16

Saving and Loading Object Selections

Save Object Selection Set

To save the list of the currently selected objects to a text file, choose Edit > Save Object Selection Set. This operation saves the names of the currently selected objects to a file that has the suffix.selset. This file is not associated with any specific project or scenario. While the file contains the names of the objects, it does not contain any object information such as location, endpoints, and attributes, model, and so forth.

The object selection set file can contain any network objects including nodes, subnets, wireless domains, links, paths, demands, and shared risk groups. However, you cannot save annotations to the selection set. Object selection sets can be referenced by Design Actions.

Load Object Selection Set

Use this menu item to read the list of objects in one or more selection set text files and then load the corresponding objects into the current scenario. You can use this operation to perform one of the following tasks:

- Add the corresponding objects to the existing set of selected objects
- Replace the existing set of selected objects with the corresponding objects

Note—If a text file contains any object names that cannot be resolved to objects in the current scenario, those object names will be listed and displayed in a status message on your screen.

Cutting, Copying and Pasting Connected Objects

The Cut, Copy and Paste commands allow you to move and copy selected objects in the Project Editor workspace. You should keep in mind that when you paste connected nodes from the clipboard, you are pasting (that is, redefining) the connecting links, demands and paths as well. (Demands and paths are used for defining traffic and are discussed in detail in the Modeling Traffic chapter.)

When a node or subnetwork is connected to another object by one or more links, paths, or demands, pasting the object is a multistep process:

1) Paste the selected node/subnetwork object(s) into the Project Editor workspace

2) Paste (redefine) any connecting links from the clipboard

3) Paste (redefine) any connecting paths from the clipboard

4) Paste (redefine) any connecting demands from the clipboard

If you cut or copy a node or subnetwork that is connected to a path, IT Sentinel copies the path to the clipboard only if the selected object is the head or tail site in the path. If you cut one or more intermediate nodes/subnets/links, IT Sentinel removes these objects from existing paths and prompts you if you want to cancel the operation. If you proceed, IT Sentinel removes only the selected sites and their connecting links from the model.

Consider the following example. In this topology, the center router (rtr_paste) is connected to its neighbors by

- Three links (to routers A, C and D)
- Two paths (path A, to router D; and path B, to routers A, B, and C)
- Two demands (from router B and to router C)

Figure 6-10 Example Topology with Links, Paths and Demands

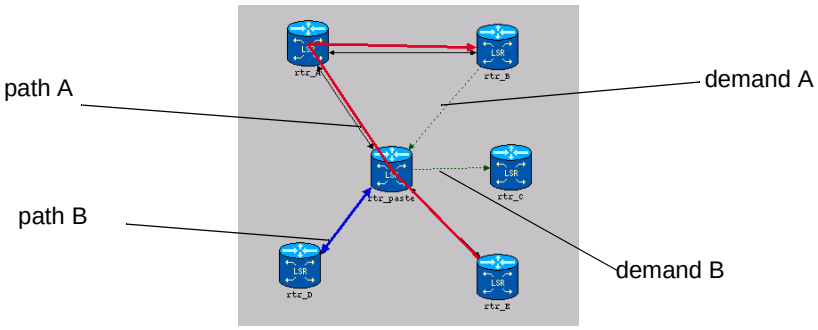

When we cut the router, it moves seven objects into the clipboard: the router, the three connecting links, path B, and the two demands. rtr_paste isn't a head or tail for path A; therefore, this path is modified rather than cut and moved to the clipboard. (A warning dialog box appears before this occurs, so you can cancel the operation if you want.)

Figure 6-11 Example Path Before and After Cut Operation

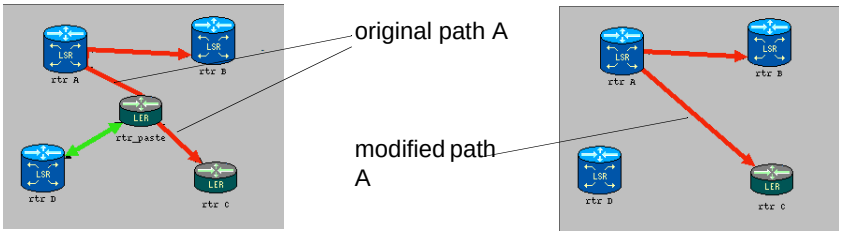

After we paste the node back into the workspace, we immediately enter link-definition mode. Note that you can check the status bar to see the current status of your paste operation. Note also that you can skip a step in the paste process, or finish the paste operation entirely, by right-clicking and choosing the appropriate option in the pop-up menu. In this example, we'll skip pasting the third link.

Figure 6-12 Paste Example: Links

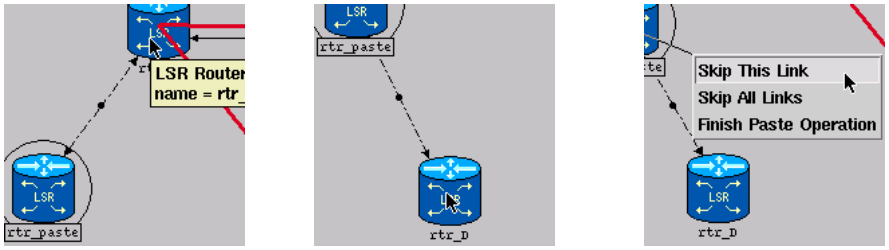

Now we automatically enter path-definition mode for path B, which was copied to the clipboard. We connect the path, right-click, and choose Finish Path Definition from the pop-up menu.

Figure 6-13 Paste Example: Paths

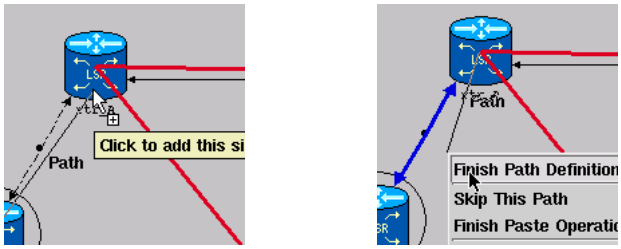

Now that we have pasted the node (step 1), links (step 2) and paths (step 3) from the clipboard, we automatically enter demand-definition mode (step 4). Because the original router had two connected flows, we now paste these demand objects. Note that both pasted demands specify the new node as the traffic source.

Figure 6-14 Paste Example: Demands

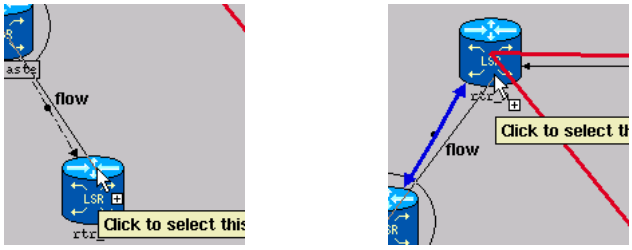

The following table summarizes the different behaviors you can expect when you cut, copy and paste connected objects.

Note—Placement parameters are in units of the subnetwork. See the subnetwork's grid to determine appropriate x, y coordinates and sizes (turn on the grid if it is not visible).

1 Click OK when you are done entering parameters.

➥ IT Sentinel draws the configuration in the workspace. It can be modified and saved like any other network model.

End of Procedure 6-16

Model Assistant

Because network models in Sentinel are usually constructed automatically as a result of imports from Device Configuration files or VNE Server, Sentinel needs a mechanism for updating the resulting imported network model to fix invalid information or provide additional information that was not present in the import source.

A model assistant file is an ASCII text file that lets you specify a set of changes to a scenario and applies the changes automatically during the import operation. You can specify multiple changes in one file, then configure the import operation to apply the file; this allows a series of changes to be performed in one step. More than one model assistant file can be applied during a Device Configuration or VNE Server import.

The Import Assistant allows you to edit certain aspects of model assistant files such as link data rates and interface connectivity, apply the data to an existing network to test the changes, and save the data so that it can be referenced when creating an import automation task step.

Procedure 6-17 Editing a Model Assistant File

1 In the Project Editor, choose Topology > Model Assistant > Edit File.

➥ The Edit Model Assistant File dialog box opens.

Figure 6-15 Edit Model Assistant File Dialog Box

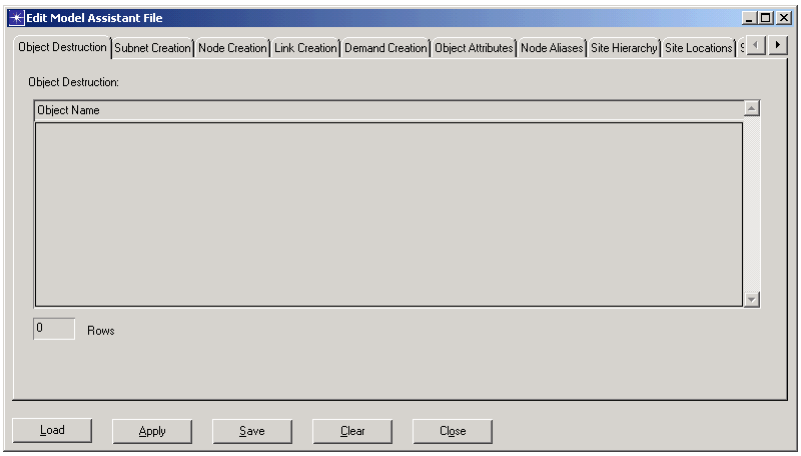

This dialog box contains multiple tabbed pages; each page has fields for entering the data you want to apply.

2 If you want to edit an existing file, click Load, choose the file to load, and click OK.

➥ The dialog box updates each tabbed page to reflect the data in the loaded file.

3 Specify the information you want to apply on each page (click the tabs to view specific pages).

- If you do not want to apply a certain type of data, leave the page empty.
- If you want to reset all settings on all pages and reenter your data, click Clear.
- If you want to save your data to a model assistant file so you can reuse it later, click Save.

4 Click Apply to apply your data to the network, or Close to close the dialog box without applying your data.

End of Procedure 6-17

Procedure 6-18 Applying Model Assistant Data to a Scenario

1 Open the project and scenario you want to update.

2 Apply the model assistant data using one of the following methods:

- If you want to apply a model assistant file without editing it, choose Topology > Model Assistant > Apply File.
- If you want to create a new file or edit an existing file, see Procedure 6-17 Editing a Model Assistant File on page ISU-6-31.

End of Procedure 6-18

Project Editor Menus

The Project Editor provides many operations for creating and working with network models. You access these operations from the Project Editor menu bar, which contains the following menus when the Project Editor is active:

- File
- Edit
- View
- Scenarios
- Topology
- Traffic
- Protocols
- Results
- Windows
- Help

In addition to the menu bar menus, two pop-up menus are frequently used when creating models:

- Workspace pop-up—contains operations related to setting the workspace view, collecting results, and viewing results.
- Object pop-up—contains operations related to setting object properties, collecting results, and viewing results.

The following tables provide a quick overview of the Project Editor menus and a reference to where each operation is discussed. The standard File, Edit, License, Windows, and Help menus are described in System Menus on page ISU-4-26.

View Menu

The View menu includes operations that affect the appearance of the editor window and its contents. See View Menu on page ISU-7-25 for descriptions of these operations.

Scenarios Menu

The Scenarios menu includes operations that provide control over the scenarios included in a project.

Table 6-4 Scenarios Menu Operations

Menu item	Description	Reference
New Scenario	Creates a new scenario within the current project.	New Scenario on page ISU-6-4
Duplicate Scenario	Creates a copy of the current scenario.	Duplicate Scenario on page ISU-6-5
Manage Scenarios	Allows you to delete, rename, and control various aspects of scenarios in the current project.	Manage Scenarios on page ISU-6-5
Previous Scenario	Switch to the next scenario in the project.	—
Next Scenario	Switch to the previous scenario in the project.	—
Switch To Scenario... >	Switches the editor window to the selected scenario.	—
Scenario Components >		
User-Defined Reports >		
Generate Report	Generate a set of output tables that list objects and attribute values of interest	User-Defined Reports on page ISU-11-3
Define Report	Define the objects and attributes to be included in an output table	Defining an Attribute Template on page ISU-9-10
Network Differences >		
Define Report	Specify the object types and attributes to consider/ignore in a network difference report.	Defining a Difference Report on page ISU-11-6
Generate Report	Generate a network difference report	Generating a Difference Report on page ISU-11-8
View Difference Data	View a network difference report	Viewing a Difference Report on page ISU-11-10
Launch Last Web Report	Open the last web report generated from a network difference report	Viewing a Difference Report on page ISU-11-10
End of Table 6-4		

Pop-Up Menus

The following sections describe the Project Editor pop-up menus used when creating models.

Workspace Pop-Up Menu

Right-clicking in any open part of the workspace brings up the Workspace pop-up menu. This menu provides access to commonly used operations for setting the workspace view, collecting results, and viewing results.

Table 6-5 Workspace Pop-Up Menu Operations

Menu item	Description	Reference
Go To Parent Subnet	Changes the view in the workspace to the contents of the next higher subnetwork.	Subnetworks on page ISU-6-18
Zoom to Selection	Zooms the view in or out so that all selected objects fill the window.	Zoom Operations on page ISU-7-22
Zoom In	Magnifies the view of the network model by 2x, centered at the cursor location.	Zoom Operations on page ISU-7-22
Zoom Out	Returns the view of the network model to the previous magnification level.	Zoom Operations on page ISU-7-22
Zoom to Window	Sets the view to include the entire subnetwork.	Zoom Operations on page ISU-7-22
Edit Selected Objects	Change the attributes on all the objects selected.	Using the Edit Selected Objects Method on page ISU-9-8
Open Device Configuration Import Log	View log for the most recent device configuration import	—
End of Table 6-5		

Object Pop-Up Menu

Right-clicking on any object brings up the Object pop-up menu. This menu provides access to commonly used operations for setting object properties, collecting results, and viewing results. The items on this menu vary depending on the type of object clicked.

The following table lists all object pop-up menu items in alphabetical order.

Table 6-6 Object Pop-Up Menu Operations (Part 1 of 2)

Menu item	Description	Reference
Advanced Edit Attributes	Opens the Attributes dialog box for the object, with advanced attributes shown.	Attributes Dialog Box on page ISU-9-1
Assign Port Connectivity		
Bring to Front	Brings the selected object to the front of its sub-layer, so it is not obscured by other objects of the same type.	Object Layering on page ISU-7-12
Edit Attributes	Opens the Attributes dialog box for the object.	Attributes Dialog Box on page ISU-9-1
Select Similar <*Objects*>	Edit the statistics (in a tabular format) on all objects that use the same underlying model.	Using the Edit Selected Objects Method on page ISU-9-8
Hide Similar <*Objects*>	Shows or hides all objects with the same “model” attribute value	—
Hide This <*Object*>	Hides the currently selected object.	—
Open Device Configuration Import Log	View log for the most recent device configuration import.	—
Open Path Viewer	Opens the Path Viewer window, which displays the currently selected path(s)	—
Select	Select the object to which the pop-up menu corresponds. This operation appears when you open the menu for an object that is not selected.	—
Select Similar <*Objects*>	Selects all objects with the same “model” attribute value.	Similar Object Selection on page ISU-9-15
Send to Back	Sends the selected object to the back of its sub-layer, where it may be obscured by other objects of the same type.	Object Layering on page ISU-7-12
Set Name	Allows you to set the name of the object. (This has the same result as editing the “Name” attribute of the object.)	—

Table 6-6 Object Pop-Up Menu Operations (Part 2 of 2)

Menu item	Description	Reference
Show Similar <*Objects*>	Shows or hides all objects with the same “model” attribute value	—
Unselect	Unselect the object to which the pop-up menu corresponds. This operation appears when you open the menu for an object that is currently selected.	—
View <*Object_type*> Description	Opens an edit pad with a description of the node or link model.	Viewing Model Information on page ISU-6-16
End of Table 6-6		

7 Viewing Networks

This chapter describes the options for visualizing, searching, and browsing your network. Among the visualization options you can specify are

- Subnetwork view properties (grid lines, distance units, etc.)
- Network maps and backgrounds
- Zooming and unzooming
- Hiding and unhiding specific objects such as demands or paths
- Customizing the appearance of link, node, and subnetwork objects

IT Sentinel include various features that are especially useful when you are working with large, complex networks:

- A network browser, which uses a treeview to organize and show network objects
- A Find Node/Link operation for finding objects of interest
- Scalable node icons, which can reduce visual clutter if a network contains tens or hundreds of nodes

Table 7-1 Viewing and Visualizing Networks: References

Topic	Reference
Network and Subnetwork Options	Subnetwork Background Properties on page ISU-7-2 Backgrounds and Maps on page ISU-7-3 Object Layering on page ISU-7-12
Visualizing Objects	Link Appearance on page ISU-7-14 Label Placement on page ISU-7-14 Icon Scaling on page ISU-7-14 Interactive Layout on page ISU-7-15 Layout Nodes Automatically (Simple) on page ISU-7-17 Hiding and Unhiding Objects on page ISU-7-13
Zoom Operations	Zoom Operations on page ISU-7-22
Working with Large Networks	Browsing and Searching Networks on page ISU-7-18 Reducing Visual Clutter on page ISU-7-24 Navigating Large Networks on page ISU-7-24
End of Table 7-1	

Subnetwork Background Properties

You can use the Set Properties operation (View > Background > Set Properties) to redefine the view properties of the subnetwork currently displayed in the Project Editor. Typically, these properties are originally set by the Startup Wizard. The following table describes the properties you can set.

Table 7-2 Background Properties Dialog Box

Property	Description
Color	Sets the color of the grid lines.
Division	Sets the number of units between each grid line.
Drawing	Sets the grid lines to be disabled (hidden), solid, or dashed.
Resolution	Adjusts the zoom level of the current subnetwork view in pixels per unit.
Subnet Layers	Sets which layers are shown in the network view and the order in which the layers appear. A layer appears in front of (and may obscure) all layers behind it. For more information about layers, see Object Layering on page ISU-7-12.
Units	Sets the scale on which the grid is based. By default, the units are the same as the parent subnetwork's. The units can be degrees, arc seconds, meters, kilometers, feet, or miles, except in the top subnetwork where units are always in degrees.
End of Table 7-2	

Setting some of the view properties automatically affects other view properties. For example, changing the *units* will also change the *resolution* and *division* properties to keep the grid lines in the same apparent location. Also, changing the resolution can cause the model to "disappear", forcing you to scroll the editor window to find it. To avoid confusion, you should normally set background properties before creating a model.

This operation does not change the overall scale of an existing network model or the defined distances between nodes. To adjust the distance between nodes after view properties have been changed, drag the objects to the new locations.

Procedure 7-1 Setting Background Properties

1 Choose View > Background > Set Properties.

➥ A dialog box opens, displaying the background properties of the current view.

2 To change units, make the desired selection from the Units pull-down menu. A subnetwork can have units of degrees only if its parent subnetwork also does.

Figure 7-1 View Properties Dialog Box

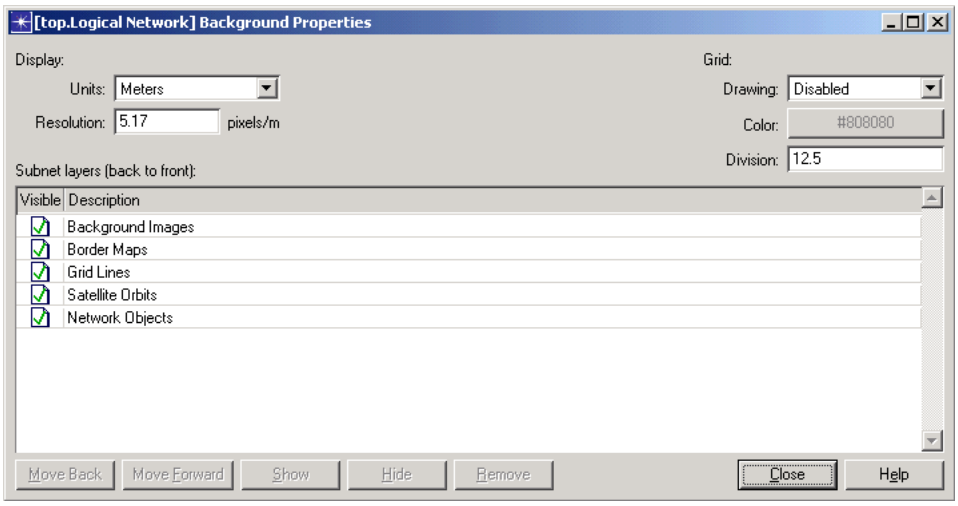

3 Set the Resolution, Division, Drawing, Color, and Subnet Layers items, as necessary.

➥ Changes made to a view property take effect immediately.

4 When you are finished, close the dialog box.

End of Procedure 7-1

Backgrounds and Maps

Backgrounds (both maps and images) represent real-world entities that are fixed in space. They provide a physical context for your network models, allowing you to interpret them more easily.

The basic difference between maps and images is that maps represent large-scale geographic entities and are described by latitude and longitude coordinates. IT Sentinel recognizes three different types of maps: border maps, image maps, and MIF maps.

Images represent small-scale entities, such as office floor plans, and are described by referencing points on an x-y axis.

Border Maps

Border maps use vector lines to define political and geographic areas such as East Asia, France, and California, using latitude-longitude coordinates to specify positions. IT Sentinel comes with a library of border maps. Border maps are visible in every subnet and cannot be edited.

Procedure 7-2 Placing a Border Map

1. Verify that your network model is set to the appropriate scale (World, Enterprise, Campus, Office, Logical, or map-specific). When you use the Startup Wizard, you can select the scale for the network. You can modify the scale of an existing model with the Set Background Properties operation (see Setting Background Properties on page ISU-6-20).
2. Display the subnet in which you want to place the map and zoom to the approximate area.
3. Choose View > Background > Set Border Map….

   ➥ The Choose Border Map dialog box appears.
4. Select the desired map.

   ➥ The map is placed in your network model.

End of Procedure 7-2

Image Maps

Image maps comprise a bit-mapped image that is tied to specific latitude-longitude coordinates. Such maps provide a more graphically detailed background than do border maps. (For background images of smaller extent, such as floor plans, see Images on page ISU-7-7.) Image maps are visible only in the subnet in which they are placed.

IT Sentinel can import two image map formats: GeoTIFF and TIFF. GeoTIFF files include latitude and longitude information, which IT Sentinel uses to place the map in the correct position. TIFF files do not include such information. When you import a TIFF file, IT Sentinel switches to map edit mode so you can enter latitude and longitude data.

Imported image map files are stored with a ***.map.i*** extension.

IT Sentinel comes with a variety of pre-imported image maps, which appear in the list presented by the Add Image Map command. The image maps were created by the U.S. Geological Survey. You can obtain additional images of this type from the Web, at http://terraserver.microsoft.com.

You can also find other maps on the Internet; these maps will often be in ***.gif*** format. You must convert such images to TIFF or GeoTIFF format before importing them.

The following procedures describe two tasks:

- How to place an already-imported image map into a network model
- How to import an image map

For information about editing an image map, see Editing Background Images on page ISU-7-10.

After they are placed in a network model, image maps can be hidden by using the Show Maps in Subnet operation (on the View > Background menu).

Procedure 7-3 Placing an Already-Imported Image Map

1. Verify that your network model is set to the appropriate scale (World, Enterprise, Campus, Office, Logical, or map-specific). When you use the Startup Wizard, you can select the scale for the network. You can modify the scale of an existing model with the Set Background Properties operation (see Setting Background Properties on page ISU-6-20).

2. Display the subnet in which you want to place the map and zoom to the approximate area.

3. From the View menu, choose Background > Add Image Map….

   ➥ The Add Map dialog box appears.

   Figure 7-2 Add Map Dialog Box

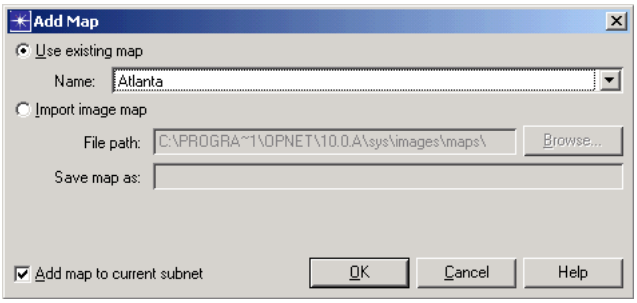

4. Verify that the Use Existing Map radio button is selected.

5. Verify that the Add Map to Current Subnet check box is selected.

6. Select the desired file name.

   ➥ The map is placed in your network model, centered in the workspace.

End of Procedure 7-3

After importing an image map, the message display may show the error message “Map is completely outside of current subnet but is added to subnet anyway.” This means that the subnet extent covers a geographic area that is different from the geographic area covered by the map you imported. You will have to import the image map into a subnet whose extent covers the coordinates of the image map for it to be visible. For example, if you import an image map of New York City while in the Paris subnet, you receive this error message and be unable to see the image map.

Procedure 7-4 Importing an Image Map

1. Verify that your network model is set to the appropriate scale (World, Enterprise, Campus, Office, Logical, or map-specific). When you use the Startup Wizard, you can select the scale for the network. You can modify the scale of an existing model with the Set Background Properties operation (see Setting Background Properties on page ISU-6-20).

2. Display the subnet where you want to place the map and zoom to the approximate area. You may want to import a border map to use as reference.

3. From the View menu, choose Background > Add Image Map….

   ➥ The Choose Map dialog box appears.

4. Click the Import radio button.

5. Verify that the Add Map to Current Subnet check box is set appropriately. You can import a map for later use even if you do not want to include it in the current scenario.

6. Do one of the following:

   - Type the path name of the file you want to import and enter a file name in the Save Map As field.
   - Click on the Browse button, select the file you want to import, and verify that the file name in the Save Map As field is acceptable.

7. Click OK.

   ➥ The file is saved in the first directory listed in your mod_dirs preference. See Preferences on page ISU-4-32 for details on modifying your mod_dirs.

   The image is placed in the workspace.

   - If the image is a GeoTIFF file, IT Sentinel uses the latitude and longitude specifications in the file to place the image map in the correct geographic position. The view may scroll so the map can be seen. A message in the message display area lists the geographic area covered by the map.
   - If the image is a TIFF file, IT Sentinel enters map edit mode and the Set Location and Size dialog box appears. See Editing Background Images on page ISU-7-10 for details on this dialog box.

At this point, you have two options: the better option is to determine and enter the correct latitude and longitude specifications for the map image. If you do this, IT Sentinel will place the map correctly next time you import it.

The other option is to exit map edit mode, import a border map, and—using the border map as a guide—place your image map in approximately the correct location.

End of Procedure 7-4

MIF Maps

MIF maps are text files of geographic information in the MapInfo Interchange File format created by MapInfo Corporation. Maps in this format are available from a variety of sources and mapping programs. The MIF specification can be downloaded from the MapInfo Corporation website at http://www.mapinfo.com.

IT Sentinel uses MIF files with a ***.mif*** extension (which contain geographic information). The directory containing the MIF files must be listed in the ”mod_dirs” preference.

MIF maps are imported as separate layers in the workspace (as described in Object Layering on page ISU-7-12). You can rearrange, hide, or remove them by choosing View > Set Background Properties.

Procedure 7-5 Importing a MIF Map

1. Display the subnet where you want to place the MIF map and zoom to the approximate area. (You may want to import a border map to use as reference.)
2. From the View menu, choose Background > Add MIF Map….

   ➥ A list of available MIF files appears.
3. Click the file name of each map you want to add to the workspace.

   ➥ Each map is place in the subnet as you click its name.

   Note—MIF data usually covers a limited area. Unless that area is visible in the subnet, you will not see imported maps until you display that area.
4. When you are finished adding maps, click Close.

End of Procedure 7-5

Images

Images comprise a bit-mapped graphic in TIFF format. They typically represent a small-scale, real-world area such as an office floor plan. (For background images of larger extent, such as a city map, see Image Maps on page ISU-7-4.) Images are visible only in the subnet in which they are placed.

When you place an image, IT Sentinel automatically switches to map edit mode so you can adjust the image position or size. Imported image files are stored with a ***.bkg.i*** extension.

The following procedures describe two tasks:

- How to place an image into a network model
- How to import an image

For information about editing an image, see Editing Background Images on page ISU-7-10.

After they are placed in a network model, images can be hidden by using the Show Maps in Subnet operation (on the View > Background menu).

Procedure 7-6 Placing an Image in a Network Model

1. Verify that your network model is set to the appropriate scale for the background image (Enterprise, Campus, or Office). When you use the Startup Wizard, you can select the scale for the network. You can modify the scale of an existing model with the Set Background Properties operation (see Setting Background Properties on page ISU-6-20).

2. From the View menu, choose Background > Add Image….

   ➥ The Add Background Image dialog box appears.

   Figure 7-3 Add Background Image Dialog Box

   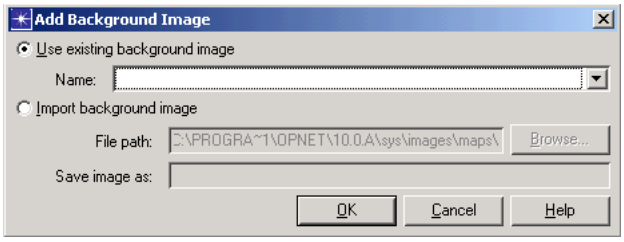

3. Verify that the Use Existing Background Image radio button is selected.

4. Select the desired background image.

   ➥ The map is placed in your network model, centered in the workspace and editing mode is turned on. See Editing a Map or Image on page ISU-7-10 for details.

End of Procedure 7-6

Procedure 7-7 Importing an Image

1 From the View menu, choose Background > Add Image…

➥ The Add Background Image dialog box appears.

2 Select the Import Background Image radio button.

3 Do one of the following:

- Type the path name of the file you want to import and enter a file name in the Save Background As field.
- Click on the Browse button, select the file you want to import, and verify that the file name in the Save Background As field is acceptable.

4 Click OK to save the file.

➥ The file is saved (with an extension ***.bkg.i***) in the first directory listed in your mod_dirs preference. See Preferences on page ISU-4-32 for details on modifying the mod_dirs preference.

IT Sentinel changes to editing mode. You can exit the editing mode by choosing View > Background > Map Edit Mode, or you can edit the background image. See Editing a Map or Image on page ISU-7-10 for details.

End of Procedure 7-7

Editing Background Images

OPNET provides a map editing mode that allows you to modify both image maps and images. In this mode, you can:

- Change the location of a map or image. “Location” means the position of the map or background image within the subnet.
- Change the size of a map or image.
- Cut, copy or paste a map or image from one scenario to another, or within the same scenario.

Procedure 7-8 Editing a Map or Image

1. From the View menu, choose Background > Map Edit Mode. In some cases, IT Sentinel changes automatically to this mode.
2. Select the map or image by clicking on it.

   ➥ Handles appear around the map or image.
3. Edit the map or image as desired, using techniques described in the following table.

Table 7-3 Edit Map or Image Operations

To...	Instruction
Cut, copy, or paste. You typically copy a map or image to a different scenario, but you can paste multiple copies in the same scenario.	Choose the desired operation from the Edit menu. When you paste a map or image to a different scenario, make sure the scale of the destination scenario matches the scale of the source scenario. You cannot, for example, paste the map of France into a scenario that is measured in feet.
Resize a map or image	Use either of two methods: • Drag one of the handles on the map. • Modify the attributes in the Set Location and Size dialog box, as described in step 4.
Move a map or image	Use either of two methods: • Drag the object to the desired position. • Modify the attributes in the Set Location and Size dialog box, as described in step 4.
End of Table 7-3	

4 To resize (or move) a map or image, right-click on it.

➥ The Map pop-up menu displays.

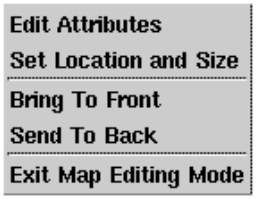

5 Choose Set Location and Size from the pop-up menu.

➥ The Set Location and Size dialog box displays. This dialog box varies, depending on whether the object to be modified is an image or a map.

Figure 7-4 Set Location and Size Dialog Box

For an image

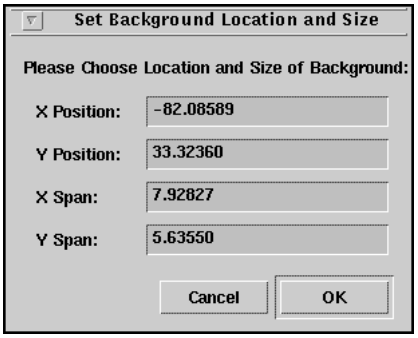

For a map

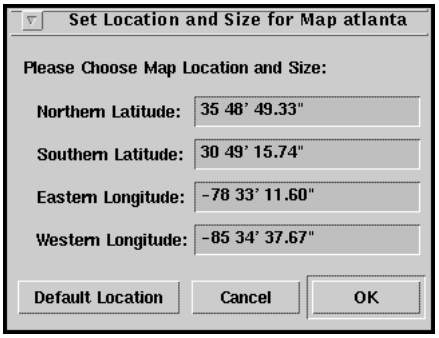

6 Follow the appropriate step:

- To resize an image—modify the “X Span” (width) and “Y Span” (height) attributes and click OK. The unit of measure is the same unit of measure as used for the subnet (degrees, kilometers, feet, and so on).
- To resize a map—modify the latitude and longitude attributes.
- To move an image—modify the “X Position” and “Y Position” attributes.
- To move a map—modify the latitude and longitude attributes.

7 From the View menu, choose Background > Map Edit Mode.

➥ The message display confirms that you have exited Map Edit mode.

End of Procedure 7-8

Object Layering

OPNET uses *layering* to help separate objects in the workspace. Each type of object (such as a node, grid line, or map) is drawn on a particular layer, with objects in higher layers overlaying—and possibly obscuring—those in lower ones. When objects of different types overlap, the frontmost object is the one that is highest in the object type hierarchy. The default hierarchy—from front to back—is as follows:

- Network objects (nodes, links, paths, and annotations)
- Grid lines
- MIF maps
- Satellite orbits
- Vector (border) maps
- Background maps and images

You can change this hierarchy with the Set Background Properties operation on the View menu.

In addition, within the network objects layer, each type of object exists in a sub-layer with other similar objects. This hierarchy is:

- Paths (see text)
- Nodes (including subnets)
- Links and paths
- Annotations

Thus, for example, links will always appear behind nodes and in front of annotations. If objects of the same type overlap, you can specify which object will appear in front by using the Bring to Front or Send to Back operation from the Object pop-up menu. These operations move the selected object within its sub-layer, relative to all other objects of its type.

Paths are treated in a special way with regard to layering, as follows:

- When created or moved to the front, a path overlays all other network objects, including nodes.
- When moved to the back, a path moves to the back of the link and path sub-layer.

Hiding and Unhiding Objects

"Show All" / "Hide All" Operations

The View menu in the Project Editor includes several operations that enable you to show or hide all objects of a given type:

- Demands (traffic flows and connection objects)
- Paths (for example, LSPs)
- Wireless domains
- Annotations created using the Annotation Palette

Hiding Individual Objects

To hide one object (demand, path, etc.), right-click on the object and choose Show this [*object*] or Hide this [*object*] from the pull-down menu.

Showing / Hiding Related Objects

You can hide all similar objects in one operation (in this case, "similar objects" means that all objects that use the same underlying model). To hide all objects of a particular type, select one or more of the objects; then choose Hide Similar [*objects*].

To unhide all objects in the Project Editor window, choose View > [*objects*] > Show All. You can also view hidden objects using the Network Browser, as shown in Network Browser on page ISU-7-18.

Viewing and Unhiding Hidden Objects

To view objects (such as demands or paths) that are currently hidden, open the Network Browser (View > Show Network Browser) and check the "Include Hidden" check box. For more information see Network Browser on page ISU-7-18.

Network Layout and Appearance

Link Appearance

A link can have several attributes (such as “color”, “thickness” and “line style”) that determines its appearance. You might want to change the attribute settings of particular links, or even derive your own link models, to make your links easier to distinguish. These are advanced attributes; to view and edit them, check “Advanced” in the Attributes dialog box.

Label Placement

The Project Editor window can move node labels automatically to reduce overlapping. If there is no space to show the node label, it is hidden unless the node is selected. To turn automatic label placement on or off, choose View > Layout > Automatic Label Placement. You can control where the label is placed with the title_autoplacing.directions preference. To allow automatic label placement to use smaller fonts, use the title_autoplacing.try_small_font preference.

Icon Scaling

The Project Editor window can resize node icons automatically to reduce overlapping and visual clutter. To turn automatic icon resizing on or off, choose View > Layout > Automatic Icon Scaling.

Because automatic icon sizing in large networks can slow down performance, autosizing is turned off if the node count reaches or exceeds the threshold specified by the icon_autosizing.large_element_count preference.

You might want to change the size of one or more site icons. For example, if an icon is 12 x 12 pixels at full size, you can scale the icon to 50% so that it appears, at most, at half-size (6 x 6 pixels).

If an icon is reduced below a certain size, the icon becomes *minimized*—that is, it appears as a simple shape rather than as a bitmap.

Figure 7-5 Scaling a Node Icon

full-sized icon | icon scaled to 50% of full size | icon scaled to 10% of full size (becomes minimized)

Procedure 7-9 Scaling One or More Node Icons

1. Select the nodes whose icons you want to scale. To scale all node icons in a scenario, left-click in the network background to ensure that no nodes are selected.

2 Choose View > Scale Selected Icons

➥ The Scale Selected Icons dialog box appears.

Figure 7-6 Scale Selected Icons Dialog Box

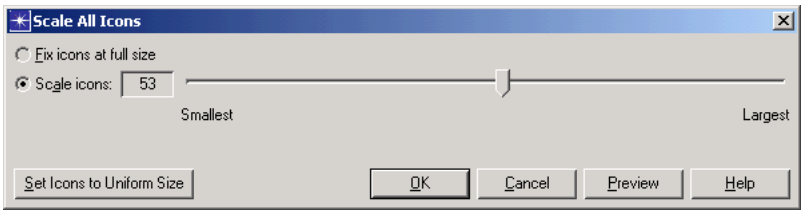

3 Scale the selected icons by the desired amount.

- To scale all icons to their full size (as defined by the icon bitmaps), select the “Fix icons at full size” radio button.
- To scale all icons to a specified percentage of their current size, select the “Scale icons” radio button; then use the slider to set the percentage.
- To scale all selected icons to the same size (that is, to the current size of the smallest icon in the selection set), click the “Set Icons to Uniform Size” button.

4 Click the Preview button to see the results of the current settings. If you don’t like the results, change the scaling options and click Preview again.

5 Click OK to apply the current settings or Cancel to revert to the original appearance.

End of Procedure 7-9

Interactive Layout

The Interactive Layout feature enables you to expand, contract, or rotate objects in a network display in order to improve visual clarity on the screen. To perform this operation, follow these steps.

Procedure 7-10 Using Interactive Layout to Adjust the Network Display

1 Choose View > Layout > Layout Nodes Interactively...

➥ The Layout Network Objects dialog box appears.

Figure 7-7 Layout Network Objects Dialog Box

2 From the Algorithm drop-down menu, select your interactive layout option (Expand/Contract is the default option).

See Table 7-4 for a description of the interactive layout options.

3 Select the objects on which you want to operate.

➥ A special green cross icon initially appears in the center of your selection set. This is the “mark” icon for the Expand/Contract and Rotate algorithms.

4 Drag the mark icon to the point on the screen where you want to anchor the selected group of objects during the layout operation.

5 For the Disperse algorithm, specify the layout region within which that algorithm will operate.

6 Click Run.

7 Repeat Steps 3 through 6, as needed.

Each iteration will add an “undoable” command to the local Undo stack. To un-do the previous layout operation, click Undo. To re-do the last Undo operation, click Redo. To put all objects in the position of their initial values (when the Layout Network Objects dialog box first opened), click Undo All.

8 Click OK.

End of Procedure 7-10

The Layout Network Objects dialog box provides the following options.

Table 7-4 Network Layout Options

Algorithm Name	Description
Expand/Contract	Moves selected objects away from or toward the green cross “mark” icon
Expand/Contract Horizontally	Moves selected objects away from or toward the mark icon horizontally
Expand/Contract Vertically	Moves selected objects away from or toward the mark icon vertically
Rotate	Rotates (by degrees) the selected objects in a counter-clockwise direction around the mark icon
Disperse	Moves selected objects within the specified layout region away from each other
End of Table 7-4	

Layout Nodes Automatically (Simple)

When you choose View > Layout > Layout Nodes Automatically (Simple), IT Sentinel automatically repositions the sites in the current subnet, trying to produce an illustrative, but not necessarily aesthetic, layout (see Figure 7-8 for an example).

Figure 7-8 Automatic Node Layout Example

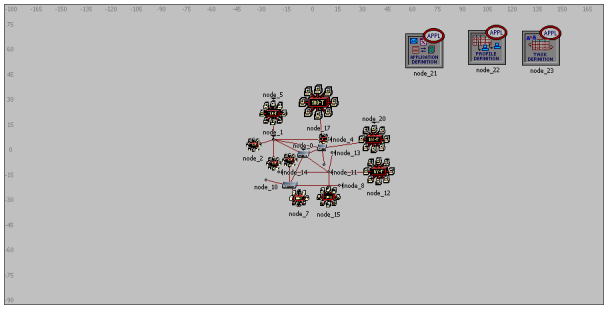

Original layout

Simple layout
(after applying the Layout Nodes Automatically (Simple) operation to the original layout)

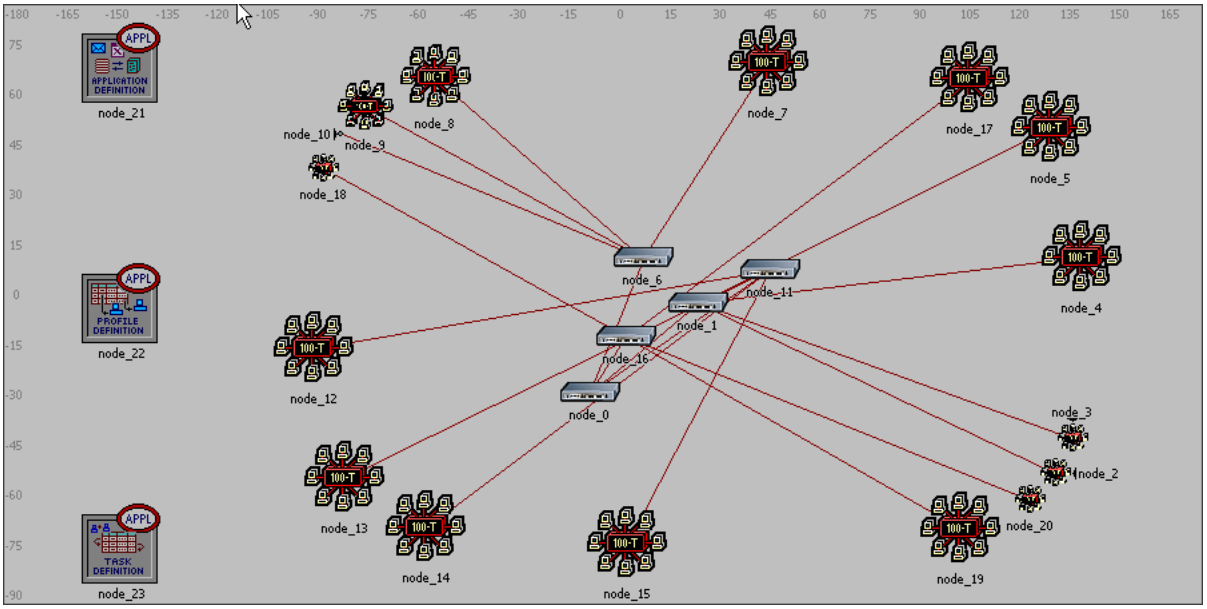

Browsing and Searching Networks

In large network models, you might have trouble finding a particular object, such as a node or subnetwork. IT Sentinel provides two features to help you in such cases: the Network Browser and the Find Node/Link operation.

Network Browser

Large, complex network models are often difficult to navigate. The Network Browser (View > Show Network Browser) uses a treeview to organize and display the objects in your network. This browser can make it much easier to find and edit objects, especially in large networks with multiple layers of subnets.

Figure 7-9 Network Browser

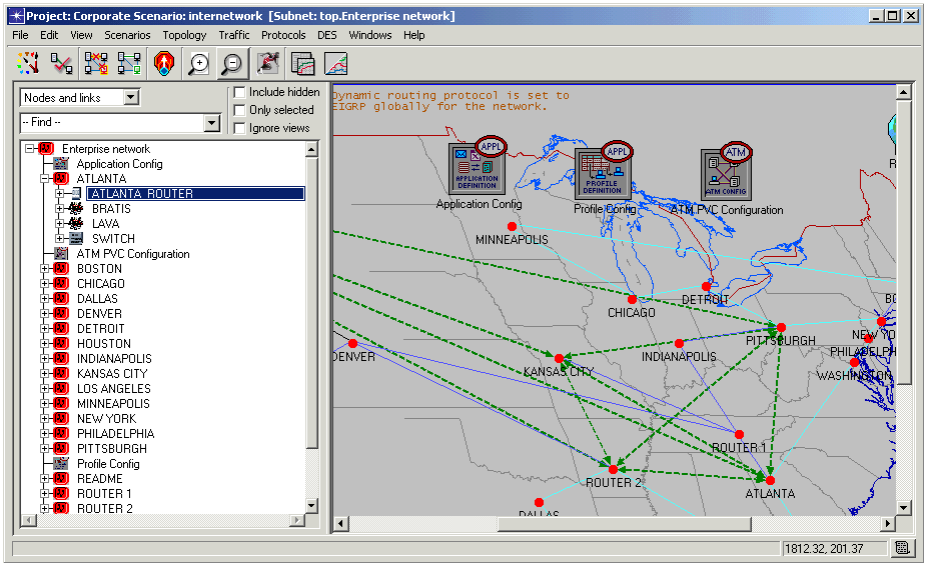

treeview pane topology pane

The browser pane and the topology pane give two different views of the same underlying network. Thus any changes you make to the network (such as selecting, editing, cutting or pasting an object) are reflected in both views.

Figure 7-10 Network Browser (Node and Link View)

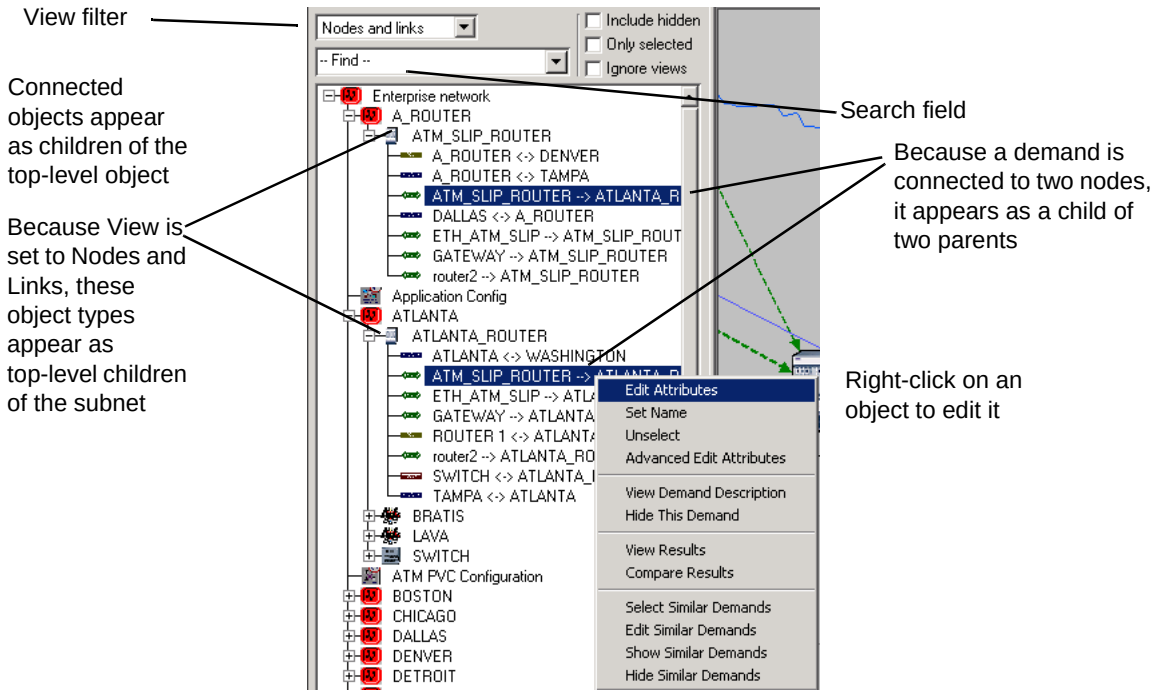

You can do the following in the network browser:

- Search for objects—To search for an object by name, enter a string in the Search field (just above the treeview pane) and press Enter.
- Select one object—To select one object, left-click on it.
- Select multiple objects—You can use several methods to select multiple objects in the browser:
  - — Hold down the Control key when you click on multiple objects.
  - — Drag the cursor to select multiple items in the treeview.
  - — Right-click on an object and choose Select Similar <*object_types*>s.
- Zoom to an object in the Project Editor—Double-click on an object in the treeview to zoom to that object in the network window.
- Edit objects—To edit one object, right-click on the object and choose Edit Attributes or Advanced Edit Attributes. To edit multiple objects of the same type, right-click on an object and choose Edit Similar <*object_types*>s

- Show or hide objects in the browser—There are several methods you can use to show or hide objects in the browser:
  - — View pull-down menu—You can use this menu to display only items of interest (such as IP routers or annotations) in the browser tree.
  - — Include Hidden checkbox—Check this box to show paths, demands and cell systems that are hidden as a result of settings in the View menu (for example, View > Demand Objects > Hide All).
  - — Ignore Views checkbox—Check this box to show objects that are excluded from the current network view.
  - — Only Selected check box—Check this box to display only those objects that are currently selected.
- Organize the treeview by object type—You set the View menu to organize the treeview by object type. These objects appear as top-level children within each subnet; any connected objects appear as children of an object. For example, suppose the View menu is set to Nodes. The treeview then displays all nodes as top-level children within each subnet. Any objects connected to a node (such as links, demands or paths) appear as children of the node.

Find Node/Link Operation

The Find Node/Link operation lets you find a node, subnetwork or link based on its name. You can specify a complete or partial name, and IT Sentinel will locate all nodes, subnets, and links that contain the string you specify. You can then jump to one of the found objects.

Procedure 7-11 Finding a Node, Subnetwork, or Link

1 Choose Find Node/Link… from the Edit menu.

➥ The Find Node/Link dialog box appears.

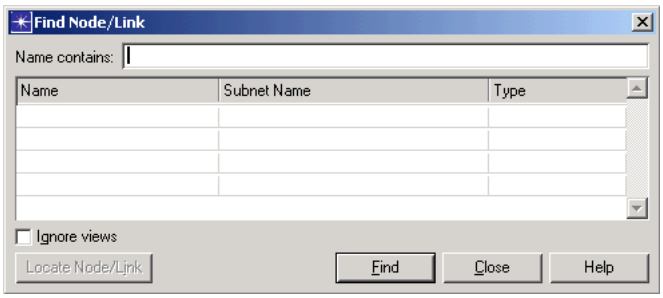

2 Type a string of characters in the Name field and click Find All.

➥ The dialog box displays a list of all nodes, subnets, and links whose name contains the string. All found objects are selected in the Project Editor window.

3 Left-click a name from the list to select it, then click Find.

➥ The Project Editor places the selected object in the center of the window. All other objects are deselected.

End of Procedure 7-11

Zoom Operations

At times, you may want to magnify the view of some portion of the network model in the workspace. This may be for display purposes or to emphasize a particular part of the network model. IT Sentinel allows you to repeatedly magnify selected areas of the workspace, as well as return to any previously selected magnification level. You can zoom in or out in several ways:

- Using the Zoom items on the View menu
- Using the Zoom toolbar buttons
- Using the Zoom items on the Workspace pop-up menu

The sections below explain the different types of Zoom operations.

Unzoom

This operation restores the previous Zoom level.

Zoom In

This operation magnifies the area displayed in the editor window. The area magnified is pre-defined: the center portion of the current display is magnified. Zoom In works with Zoom Out, which restores the area displayed before the Zoom In operation took place.

Note—Zoom In is not the same as Zoom to Rectangle. Zoom to Rectangle allows you to define the area magnified.

Zoom Out

This operation displays a larger area of the model in the editor window.

If you magnified a region by choosing the Zoom In command from the workspace pop-up menu or the View menu, Zoom Out restores the view to the exact previous zoom level. If you chose the workspace Zoom In command multiple times, each operation of the Zoom Out command restores the previous zoom level.

If you magnified a region by choosing the Zoom to Rectangle command from the View menu, and then defined an arbitrary region, you may need to choose Zoom Out several times before the magnification approximates the original zoom level.

Note—Zoom commands on the View and workspace pop-up menus affect the network view, even if an analysis panel is the active window.

Zoom to Rectangle

This operation allows you to get a close-up view of a model at a specified location within the Project Editor. This operation may be applied multiple times to provide an ever-closer view.

If you choose Zoom to Rectangle from the View menu or click on the Zoom button, the cursor changes to a bar and you drag to define the rectangular-shaped region to be magnified. If you choose Zoom In from the View menu or the Workspace pop-up menu, IT Sentinel zooms to a pre-defined area.

Zoom to Selection

This operation zooms the view in or out so that all selected objects fill the editor window.

Zoom to Window

This operation displays all objects in the current network or subnetwork. If a subnet's extent is very large, Zoom to Window may cause all objects to be clustered together at one point on the screen, making the network difficult to view. If this happens, select all objects in the cluster and use Zoom to Selection.

Note—Zoom to Window clears all previous views in the Project Editor, and cannot be undone by using Zoom Out. (This simply shrinks the subnetwork view in the editor window).

Working with Large Networks

This section outlines techniques for viewing and navigating very large networks. If a particular scenario contains a lot of nodes and/or subnetworks, you might encounter the following problems:

- Visual clutter—If a particular subnetwork contains a lot of nodes, you might find that the objects are grouped too closely to distinguish easily.
- Navigation—If your network contains numerous nodes and subnets, you might have difficulty finding particular objects of interest.

This section provides some pointers for working with large networks and provides references for further information.

Reducing Visual Clutter

The following operations can help you reduce visual clutter in your network:

- Hide link arrowheads—You can change the thickness, line style, and arrowhead settings for individual links (right-click on a selected link and choose Advanced Edit Attributes).
- Scale network icons—One effective way to minimize visual clutter is to reduce the size of the node/subnet icons in your scenario. For more information, see Icon Scaling on page ISU-7-14.

Navigating Large Networks

You might want to use the Network Browser (View > Show Network Browser) if your network contains many objects or subnetworks. You can also use the Find Node / Link operation (Edit > Find Node/Link) to search for objects of interest. For more information, see Browsing and Searching Networks on page ISU-7-18.

View Menu

The View menu includes operations that affect the appearance of the editor window and its contents. (See Project Editor Menus on page ISU-6-33 for information on other menu operations.)

Table 7-5 View Menu Operations (Part 1 of 3)

Menu item	Description	Reference
Show Network Browser	Displays the hierarchical structure of all subnets, nodes, and links in the current network model and allows you to set their attributes.	Network Browser on page ISU-7-18
Background >		
Set Properties…	Sets the properties of a subnetwork's coordinate system, including units, resolution, and grid display.	Setting Background Properties on page ISU-6-20
Set Border Map...	Imports an outline map of a geographic area into the workspace, providing a graphic frame of reference for a network.	Border Maps on page ISU-7-4
Add Image Map...	Imports a bit-mapped graphic of a geographic area into the workspace, providing a graphic frame of reference for a network.	Image Maps on page ISU-7-4
Add MIF Map…	Imports a map in MIF format into the workspace. MIF is a common format for maps of roads, political boundaries, and other data.	MIF Maps on page ISU-7-7
Add Image...	Imports a bit-mapped image of a small area such as a floor plan into the workspace, providing a graphic frame of reference for a network.	Images on page ISU-7-7
Map Edit Mode	Toggles map edit mode for image maps and background images. In this mode, you can: • Cut, copy, or paste • Modify the size • Modify the location	Editing Background Images on page ISU-7-10
Show Maps in Subnet	Toggles display of image maps and background images.	—
Refresh Workspace	Redraws the display to accurately portray the window and model.	—

Table 7-5 View Menu Operations (Part 2 of 3)

Menu item	Description	Reference
Layout >		
Automatic Icon Scaling	Turn automatic icon scaling on or off.	Icon Scaling on page ISU-7-14
Automatic Label Placement	Turn automatic label placement on or off	Label Placement on page ISU-7-14
Scale Node Icons Interactively...	Set the icon size for nodes, subnets, and utility objects.	Icon Scaling on page ISU-7-14
Layout Nodes Interactively...	Change the layout of nodes in the Project Editor workspace to improve visual clarity on the screen	Interactive Layout on page ISU-7-15
Layout Nodes Automatically (Simple)	Apply a simple, but fast algorithm to produce a “reasonable” network layout	Layout Nodes Automatically (Simple) on page ISU-7-17
Zoom >		
To Rectangle	Allows you to drag a rectangle that demarcates the region to display.	Zoom to Rectangle on page ISU-7-23
To Selection	Zooms the view so all selected objects are displayed.	Zoom to Selection on page ISU-7-23
To Window	Displays all objects in the current network or subnetwork.	Zoom to Window on page ISU-7-23
Unzoom	Restores the previous zoom level.	Unzoom on page ISU-7-22
Subnets >		
Go To Parent Subnet	Changes the view in the workspace to the contents of the next higher subnetwork.	—
Expand Selected Subnets	Displays the contents of all selected subnets and allows certain operations to be performed on the displayed objects.	Expanding a Subnet on page ISU-6-20
Collapse Selected Subnets	Restores selected subnets to their typical collapsed state, where they are represented by one icon.	Collapsing a Subnet on page ISU-6-21
Demands >		
Show All	Displays all demand objects in the Project Editor.	—
Hide All	Hides all demand objects in the Project Editor.	—
Paths >		
Show All	Displays all path objects in the Project Editor.	—
Hide All	Hides all path objects in the Project Editor.	—

Table 7-5 View Menu Operations (Part 3 of 3)

Menu item	Description	Reference
Annotations		
Show in subnet	Toggles display of annotations. The setting you specify is on a per subnet basis. You can hide annotation in one subnet and show it in another. The setting is saved with the model.	Object Palettes on page ISU-6-8
Visualize Protocol Configuration >		
IP Interface Status	Shows whether the connected interface is active or shutdown.	—
IP Routing Domains	Shows the routing protocols that are configured on the router interfaces.	—
IP QoS Configuration	Shows the IP QoS configuration on the links.	—
IP Security Configuration	Shows the packet filtering security configured on the links.	—
IP Tunnel Configuration	Adds colored path objects to the workspace to represent the IP tunnels configured in the network.	—
BGP Peers	Shows the BGP peering (neighbor) relationships in the network.	—
OSPF Area Configuration…	Shows the OSPF areas configured on the router interfaces.	—
ATM Routing Domains	Shows which ATM routing protocols that are configured on the switch ports.	—
VLAN Configuration…	Shows the VLANs configured in the network.	—
Clear Visualization	Removes all protocol configuration visualizations from the workspace.	—
Visualize Network Configuration >		
End of Table 7-5		

8 Modeling Traffic

After you create a network topology (manually or by import), the next step is to add traffic to the network. You can create traffic in your network manually by setting attributes on various network objects.

In IT Sentinel, you can model background traffic.

“Background traffic” is analytically modeled traffic that affects the performance of explicit traffic by introducing additional delays. Unlike explicit traffic, background traffic can affect not only discrete event simulations, but also flow analyses. Discrete event simulations that include background traffic use the hybrid simulation model. This model includes the effects of background traffic to calculate queue build-ups on intermediate devices and delays based on the queue length, at any time during in the simulation. Because each packet that produces traffic on the network is not explicitly modeled, using background traffic can speed up simulations considerably.

Background traffic takes three forms in IT Sentinel:

- Traffic Flows—A traffic flow describes an end-to-end flow of traffic from a source to one or more destination nodes. You can create traffic flows manually, using traffic flow objects (covered later in this chapter).
- Device/Link Loads—This type of traffic (also called “static background utilization”) represents traffic as a background load on a link or node object. Unlike a traffic flow, which can span multiple links and nodes, a traffic load is “static” and applies to one object. You can convert existing link loads to traffic flows, which allows flow analyses to account for these loads.
- Application demands—You can use application demands to represent background traffic flowing between two nodes. Besides background traffic, you can also configure application demands to be purely discrete (explicit) traffic, or a combination of these two (hybrid).

Traffic References

The following table contains references to information on creating and working with network traffic.

Table 8-1 Network Traffic: References

Topic	Reference
Methodology for Representing Network Traffic	See the Creating a Unified View of the Network methodology
Background Traffic (overview and detailed discussion)	Working with Background Traffic on page ISU-8-19 Background Traffic and Statistics on page ISU-8-29 Limitations of the Background Traffic System on page ISU-8-31
Traffic Flows	Traffic Flows on page ISU-8-3
Device/Link Loads	Device/Link Loads on page ISU-8-16
Paths	Using Paths to Represent Virtual Circuits on page ISU-8-12
End of Table 8-1	

Traffic Flows

Creating Traffic Flows and Profiles

Traffic Flows and Demands

You create traffic flows in your network using *traffic flow* objects. A traffic flow is a network object that connects two nodes (like a link object) and specifies:

- A traffic source and a traffic destination
- A period of simulation time, which can be divided into time slots
- The rate of traffic (in bits-per-second, packets-per-second, or some other measure) from source to destination during each time slot

Traffic flows are a type of *demand object*. Demands are used to model potentially variable relationships between nodes. Traffic flows are useful for defining IP and ATM end-to-end traffic; therefore, this section focuses on this type of demand.

You can browse, view and edit the traffic flows in your network using the Flows Browser, which is described in Viewing Flows and Connections on page ISU-8-10.

Traffic Profiles

Every flow's attribute list includes one or more *traffic profiles* that specify the traffic intensity for that flow object. "Traffic Intensity," in this case, means the rate of source-to-destination traffic for a given simulation time period. A traffic profile stores a series of x-y value pairs, in which

- X values represent time slots, in seconds
- Y values represent rates of traffic (in bits/second, packets/second, or some other measure) for each corresponding time slot

When you view or edit a profile attribute, OPNET displays that attribute's values using both a table and an x-y graph, as shown in the following diagram.

Figure 8-1 Traffic Profile Attribute for a Flow Object

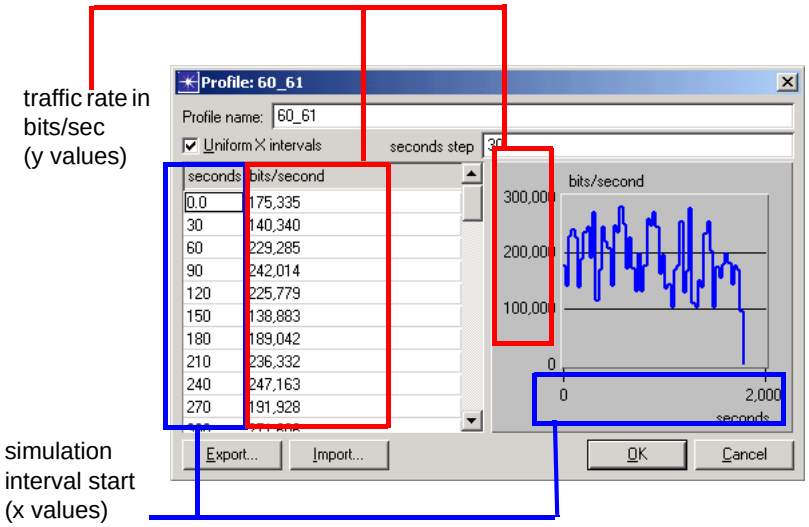

Note that a flow specifies traffic from a source to a destination; it may also include a *path object* to specify that the traffic takes a certain route. If a flow does not include a path, the routing protocols used in the network determine the routes used by the traffic. (Path objects are covered in detail in Using Paths to Represent Virtual Circuits on page ISU-8-12.)

The following diagram shows a network with all three types of connection objects: links (solid line/two arrowheads), flows (blue dashed arrows), and paths (solid green line/one arrowhead indicates direction of path).

Figure 8-2 Example Network with Links, Demands and Paths

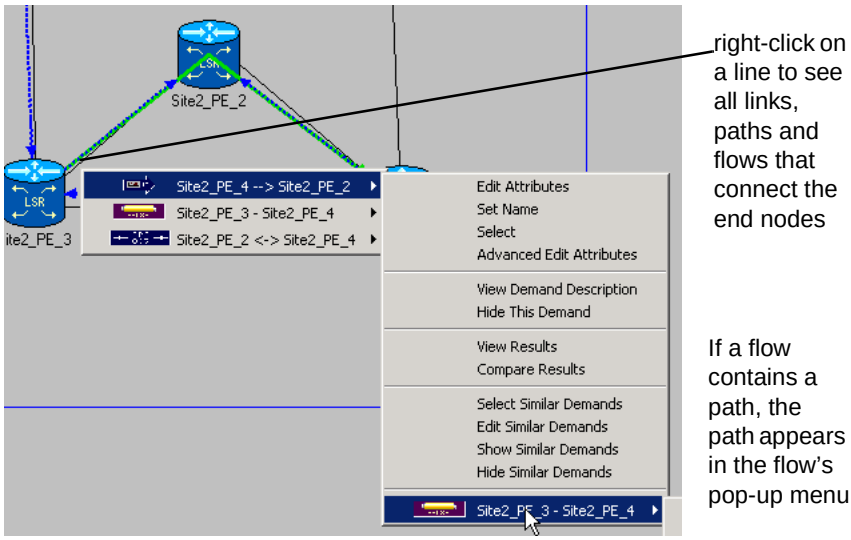

When you right-click on a traffic flow, the pop-up menu operations see “demands” rather than “flows.” A traffic flow is a type of demand object; when you right-click on a traffic flow and choose “Show Similar Demands,” for example, IT Sentinel shows all flow objects that use the same underlying model.

Demand Model Options

Each flow model has a set of options that determine the type of flow object you can create. You can view model options in the Demand Model Description dialog box, which you open by right-clicking on a model in the object palette.

Figure 8-3 Demand Model Description Dialog Box

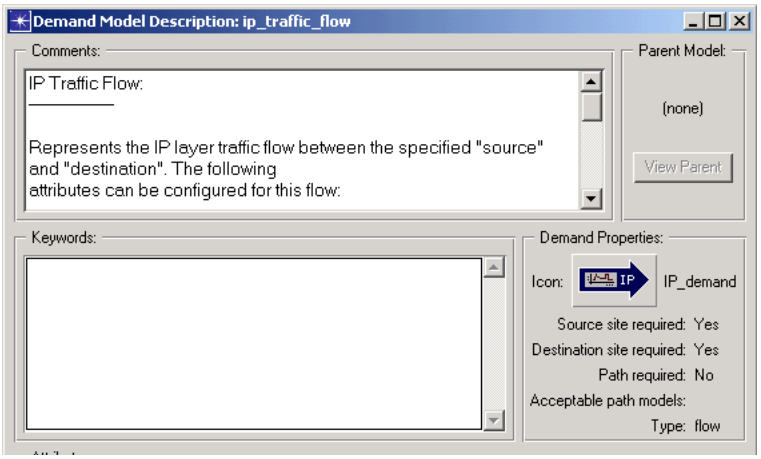

Before you create a flow, verify that the underlying model options support the kind of object you want to create. The three most important considerations are the demand type (connection or flow), the supported attributes, and whether the demand requires an associated path object.

Demand Types: Connections and Flows

There are two types of demand objects: connections and flows.

- Connections are useful for defining pipelines or reservations between nodes such as PVCs or reserved bandwidths.
- Flows are useful for defining end-to-end background traffic. A traffic flow generally defines two nodes (traffic source and destination) and one or more traffic profiles that specify the rate of traffic.

In most cases flows, not connections, are a better choice to model traffic flows.

Traffic Profiles

Before you create a demand object, check the attribute list of the model you plan to use to ensure that the model has the correct profile type. For example, if you want to specify traffic in bits-per-second, be sure to choose a model that includes that attribute.

Point Demands

Some demands, typically those used for multicast traffic, support only a source or a destination, but not both. You can recognize such point demands in the Demand Model Description dialog box because only one of the “Source site required” and “Destination site required” properties will be set to Yes.

Demands and Paths

If a demand is not associated with a path object, the routing protocols in your network determine the path that the traffic takes.

If you want the traffic to take a specific route, you can have a path refer to a path object. In some cases, this is required: the “Path required” model option determines whether a demand requires an associated path. A demand may also be restricted to certain types of paths, as specified by the “Acceptable path models” option.

Creating Traffic Flows

Creating a traffic flow is a two-step process:

- Create a flow object that connects the source and destination nodes, as described in Creating a Traffic Flow on page ISU-8-6.
- Specify the amount of traffic intensity for the flow, as specified in Specifying Traffic for a Flow on page ISU-8-7.

Creating a Traffic Flow

Procedure 8-1 Creating a Traffic Flow Object

1. If necessary, open an object palette and configure it with the flow models you want.
2. Select a model in the object palette by clicking on it. (See Creating Traffic Flows and Profiles on page ISU-8-3 for more information on the various demand types available).

   ➥ The cursor changes appearance to indicate that you are in flow-creation mode.

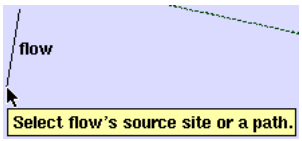

3. Click on the *source node* of the traffic flow.
4. If you want to include a path object in your flow definition, click on the path.
5. Click on the *destination node* of the traffic flow.

   ➥ A new traffic flow is created. By default, OPNET uses a blue dotted arrow to represent the traffic direction of the traffic.

6 (*Required-path flows only*): If there is no path in your flow, you must include one now. Click on the path object to complete the flow definition.

Note—When you finish defining a flow, you are still in flow-definition mode. You can create another flow (based on the same underlying model) by repeating the previous steps. To exit flow-definition mode, right-click and choose Abort Flow Definition from the pop-up menu.

After you define the flow, you can specify the traffic levels associated with the flow object, as described in Specifying Traffic for a Flow on page ISU-8-7.

End of Procedure 8-1

After you define the flow, you can specify the traffic levels associated with the flow object, as described in Specifying Traffic for a Flow.

Specifying Traffic for a Flow

You specify traffic levels for a traffic flow by editing the traffic attributes of the associated traffic flow object (see Creating Traffic Flows and Profiles on page ISU-8-3 for more information). There are several methods available to do this:

- Specify traffic levels manually—You can enter data manually into the traffic attributes of the flow (see Manually Specifying Traffic on page ISU-8-7)
- Assign an existing traffic profile—You may want to assign an existing traffic profile to a flow (see Assigning a Traffic Profile on page ISU-8-9)
- Import data from an external file—You can import traffic data from an external text (*.txt) or comma-separated value(*.csv) file directly into a profile (see Importing Traffic Data Into a Profile on page ISU-8-9)

Manually Specifying Traffic

Procedure 8-2 Specifying Traffic Levels on a Flow Object

1 Right-click on the flow and choose Edit Attributes to open the Attributes dialog box for the flow object.

2 Click in the Value field of the appropriate traffic attribute and choose Edit from the pull-down menu.

➥ The Attribute Profile dialog box appears.

To enter a new value, click in the first empty row in the left (time) column

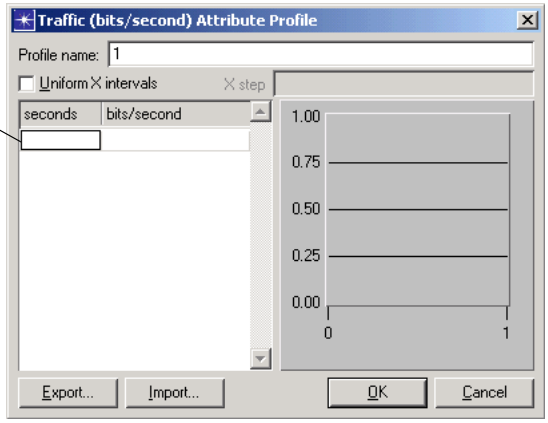

3 Specify a name for the profile, if desired, in the Name field.

4 Specify a uniform time step for the x axis (simulation time), if desired. The Uniform X Intervals check box and Time Step field allow you to specify uniform time intervals for the profile.

5 Enter the x/y value pairs by clicking in the x-value (left) and y-value (right) columns, respectively.

➥ The graph on the right automatically updates to reflect the value pairs you enter.

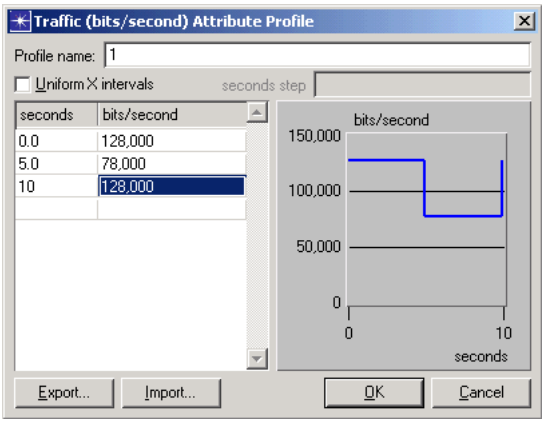

If you specified a uniform time step (that is, if the Uniform X Intervals check box is enabled), you can enter only x values that are multiples of the time step. To specify a value that is not a multiple of the time step, disable the Uniform X Intervals check box.

6 Click OK to close the Profile dialog box, then OK to close the Edit Attributes dialog box.

End of Procedure 8-2

Assigning a Traffic Profile

Procedure 8-3 Assigning an Existing Traffic Profile to a Flow

1. Right-click on the flow and choose Edit Attributes to open the Attributes dialog box.
2. Click in the Value field of the traffic attribute you want to edit and choose Select from the pull-down menu.

   ➥ The Attribute Profile dialog box appears.
3. Choose the library that contains the profile you want to use in its Profile Library list.

   ➥ The traffic profiles contained in the selected library appear in the Profile list.

   Note—If the profile library you want to use is not in the list, the most likely reason is that the file's directory is not listed in your model directories. Choose File > Add Model Directory and specify the directory as a Source Directory.
4. Choose a profile from the Profile list and click OK; then click OK to close the Edit Attributes dialog box.

End of Procedure 8-3

Importing Traffic Data Into a Profile

You can import traffic data into a profile from either a text (.txt) or a comma-separated value (.csv) file.

Import Requirements

A data file for a traffic profile should describe one traffic flow only. You should ensure that the traffic data measured fits the attribute. For example, if you are importing into an attribute called "Flow Intensity (bits/second)", make sure that the data you want to import describes traffic in bits per second, not packets per second or some other measure.

The following diagram shows the required file format for your traffic data. The x and y header fields must match the headers of the target profile.

Figure 8-4 Required Data File Format

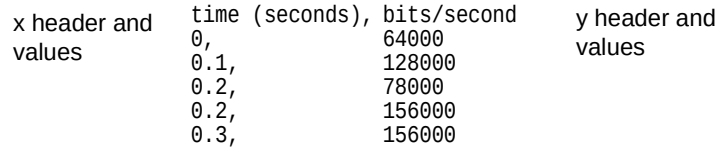

Procedure 8-4 Importing Traffic Data into a Profile

1. To open the Edit Attributes dialog box for a flow, right-click on the flow and choose Edit Attributes.

2 Click in the Value field of the traffic attribute you want to import and choose Select from the pull-down menu.

➥ The Attribute Profile dialog box appears.

3 Click the Import button and choose the data file you want to import.

4 Click OK to close the Profile dialog box, then OK to close the Edit Attributes dialog box.

End of Procedure 8-4

Viewing Flows and Connections

IT Sentinel includes a Demand Browser that enables you to navigate and edit all flow or connection objects in a scenario. To open this browser, choose Topology > Open Connections Browser or Traffic > Open Flows Browser. These two browsers show different types of demand objects but otherwise have the same interface.

Figure 8-5 Traffic Flows in the Flows Browser

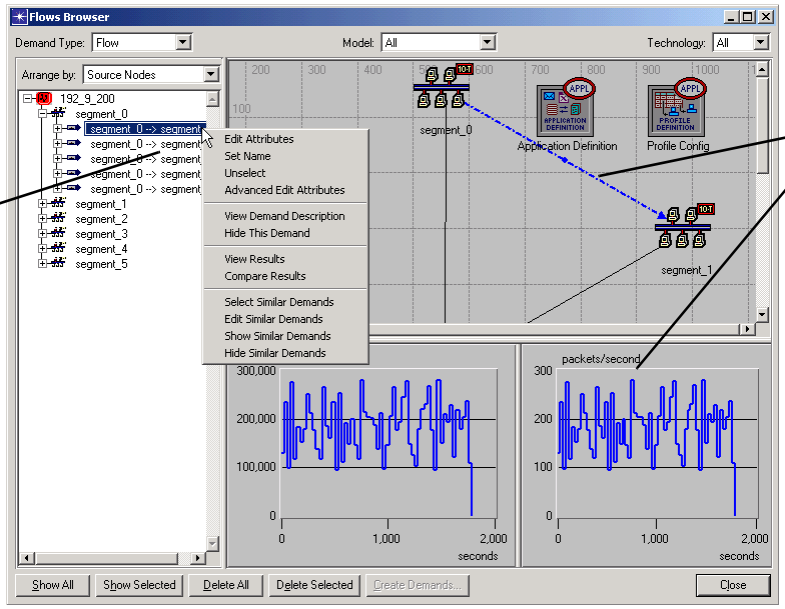

Right-click on a demand in the treeview or network window to open the Demand pop-up menu

Network and profile panes show location and traffic intensity of selected demand

To edit a profile or assign a different profile to the demand, right-click in a profile panel

You can do the following operations in the Flow and Connection Browsers:

- Find and view specific objects—The Demand Type, Model and Technology pull-down menus determine the types of demand objects that appear in the treeview. When you select an object in the treeview, the browser window on the right displays the selected object's location and traffic profiles (if appropriate).
- Arrange the treeview by type—Set the Arrange By pull-down menu (above the treeview) to arrange the treeview by object type: traffic flows, demand sources/destinations, etc.

- Edit a specific object—To open a context-sensitive menu for a selected object, right-click on it (in the treeview or in the network pane) and choose Edit Name, Edit Attributes, or Advanced Edit Attributes.
- Edit multiple demands—To edit the selected demand or other demands based on the same model, right-click on a selected demand and choose Edit Attributes, Advanced Edit Attributes or Edit Similar Demands from the pull-down menu.
- Show or hide demands—To show or hide all demands of a particular type, select a demand, right-click and choose Show Similar Demands or Hide Similar Demands.

  To show all demands, click the Show All button at the bottom of the browser window. To show only demands that are currently selected (and hide all others), click the Show Selected button.
- Delete selected demands or all demands—Click the Delete All or Delete Selected button.

Note—To undo a delete-demand operation, choose Edit > Undo in the Project Editor window.

- Edit a traffic profile or assign a different profile to a demand—Right-click in either of the profile panels (lower right) and choose Select Profile (to select a new profile for the demand) or Edit Profile (to edit the existing profile).
- View traffic for a specific demand if multiple demands are selected—When you select multiple demands in the browser, the profile panels (lower right) are initially blank. To show the profiles for a specific demand, left-click in a profile panel and choose the demand in the pop-up menu.

Demand Pop-Up Menu Operations

Right-click on any demand object to open the Demand pop-up menu. For information on the operations available on this menu, see Table 6-6 Object Pop-Up Menu Operations on page ISU-6-36.

Using Paths to Represent Virtual Circuits

You can create *path objects* between nodes or subnets. Any protocol model that can use logical connections or virtual circuits (such as MPLS) can use paths to represent traffic routes.

Paths have no built-in simulation behavior: the underlying protocol models determine when and how to use a particular path. See the relevant model documentation for information about how particular models use paths.

A path is a defined route between two *sites*; a site can be either a fixed node object or subnetwork. The process of defining a path is similar—though not identical—to defining a link. First, configure an object palette to include the path models you want (as described in Object Palettes on page ISU-6-8); then select a model and click on objects in the Project Editor workspace to connect them.

Path Model Options

Each path model has a set of options that determine the path connectivity requirements and the types of objects that you can include in a path definition. You can view a model's options in the Path Model Description dialog box, which you open by right-clicking on a model in the object palette. You should verify that the underlying model's options allow you to create the type of path you want. If you cannot find a path model whose options exactly match the type of path you want to create, you may want to derive a new model from an existing one.

Figure 8-6 Path Model Description Dialog Box

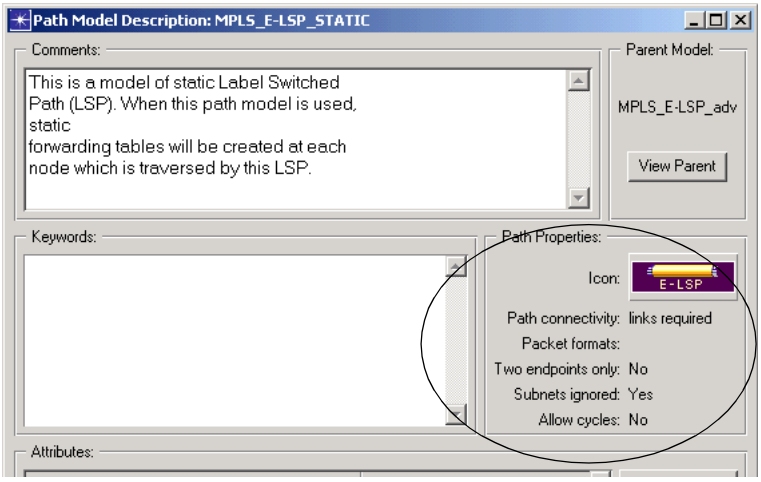

path options

Path Connectivity

A path object defines, at the very least, two end sites. It may also define a partial or complete path (that is, a set of connecting sites and links) between these sites. The “Path Connectivity” option defines the extent to which you can define the intermediate path. The three possible settings are:

- Links Required—In a required-link path, you must specify a simplex or duplex point-to-point link between every site in the path; all such links are included in the path definition.
- Links Ignored—In an ignored-link path, you can only include sites in the path definition. The underlying models fully determine routes between the included sites.
- Links Optional—This is a combination of the required-link and ignored-link types. You must define the head, tail and any intermediate sites, defining any connecting links between sites is optional.

The following diagram shows a router network that connects a site in Canada to a site in Texas. Traffic constrained by the required-link path (left) always follows the exact route defined between Canada and Texas. If the ignored-link path (center) defines a route, the underlying models fully route any traffic using this path. The optional-link path (right) defines a partial route that includes the link between Canada and router E; the underlying models route the traffic between router E and Texas.

Keep in mind that in a situation like this, with multiple paths defined between the same two sites, the underlying models choose the actual path taken by a packet.

Figure 8-7 Multiple Paths Between Two Sites

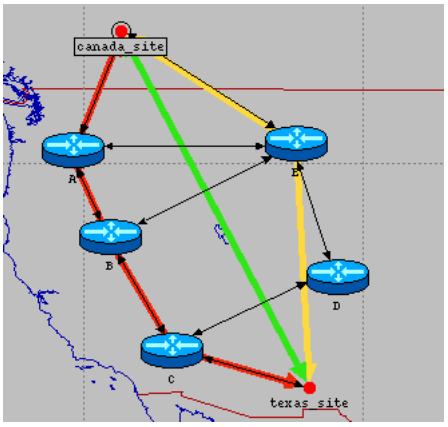

Other Path Options

Other path options are:

- Packet Formats—Each path has a defined set of supported packet formats. You cannot define a link as part of a path unless that link supports all formats included in the path model's Packet Formats field. (Note that if a path supports no packet formats, you can include any link in the path definition.)
- Two Endpoints Only—If this option is set to Yes, you can create single-segment paths only.
- Subnets Ignored—If this option is set to Yes, you cannot include a subnetwork object in a path definition. (However, you can still include nodes from different subnets in the same path.)
- Allow Cycles—If this option is set to No, you can only include each link or site once in a path definition.

Creating a Path Object

Procedure 8-5 Creating a Path Between Two Sites

1. If necessary, open an object palette and configure it to include the path models you want.
2. Left-click on a path model in the object palette to select it.

   ➥ The cursor changes appearance to indicate that you are in path-creation mode.

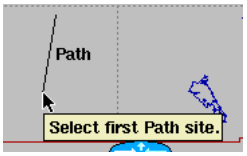

A path object's underlying options (described in Path Model Options on page ISU-8-12) determine the objects you can include in your path definition. In path-creation mode, IT Sentinel provides extensive visual feedback to validate additions to your path. If you move the cursor over an invalid link or site, the cursor changes appearance and a tooltip appears to indicate why that choice is invalid.

Figure 8-8 Visual Feedback in Path-Creation Mode

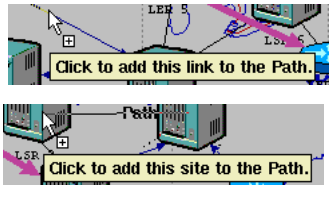

Feedback for valid objects

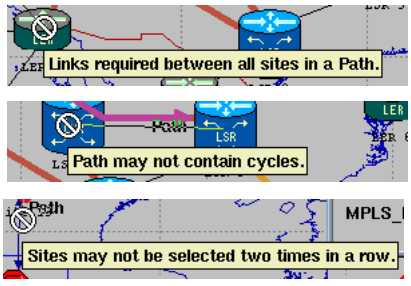

Feedback for invalid objects

3 Left-click on the site you want to define as the head of your path.

4 (Required-link and optional-link paths only): Left-click on any sites and links that you want to include in your path.

You must click on a link to include it in the path definition if

- You are defining an links-optional path;
- You are defining a links-required path and want to include two sites that have multiple links connecting them.

5 Repeat step 4 until you have clicked on the tail site and the path definition is complete.

6 Right-click in the project workspace and choose Finish Path Definition from the pop-up menu. IT Sentinel auto-assigns a unique name to the new path.

Note—When you finish defining a path, you are still in path-definition mode. You can create another path (based on the same underlying model) by repeating steps 3 through 7. To exit path-definition node, right-click and choose Abort Path Definition from the pop-up menu.

End of Procedure 8-5

Path Pop-Up Menu Operations

Right-click on any demand object to open the Path pop-up menu. For information on the operations available on this menu, see Table 6-6 Object Pop-Up Menu Operations on page ISU-6-36.

Device/Link Loads

A *Device/link load* is a static throughput on a device or link. Unlike a traffic flow (which flows into, out of, and through multiple objects), a device/link load applies to one object only. Like traffic flows, device/link loads can make discrete event simulations run faster and consume less memory than simulations that model all traffic explicitly.

If a link or node model supports a background load, it has an attribute that specifies the traffic during a specific interval of simulation time. Node and link models have different attributes and define background loads differently:

- A device load is usually defined as *a percentage of the total capacity* on a server, LAN, or other node. Node models have an attribute called "Background Utilization," which specifies the load as a percentage of the total utilization on that node. For example, a LAN model might contain an attribute called "LAN Background Utilization," which specifies the baseline load on the modeled LAN.
- A link load is usually defined as *an absolute level* of traffic in bits/second using a profile attribute type (described in Traffic Profiles on page ISU-8-5). Link models have an attribute called "Background Load," which defines the traffic throughput in bits/second. You can define two unidirectional loads (one in each direction); a link model might have additional attributes that define the traffic loads in more detail.

Most standard models have a default load level of 0 (0 bits/second for a link and 0% utilization for a node).

Specifying Device/Link Loads Manually

To specify device/link loads manually, you must configure the attributes of individual objects.

Procedure 8-6 Creating Link Loads Manually

1 Right-click on a link and choose Edit Attributes.

➥ The Attributes dialog box opens.

2 Expand the tree for the “Background Load” attribute. The “Background Load” is a compound attribute that defines two unidirectional loads (“Node A > Node B” and “Node B < Node A”).

Background Load	(...)
Average Packet Size (A -> B) (bytes)	NONE
Average Packet Size (B -> A) (bytes)	NONE
A -> B	NONE
B -> A	NONE

For each load, do the following steps.

2.1 To specify a unidirectional traffic level, click in the Value field of the “A -> B” or “B -> A” attribute and click one of the following:

- Edit (to specify the traffic manually)
- Select (to use a profile from an existing library).

Figure 8-9 Traffic Profile for a Link Load

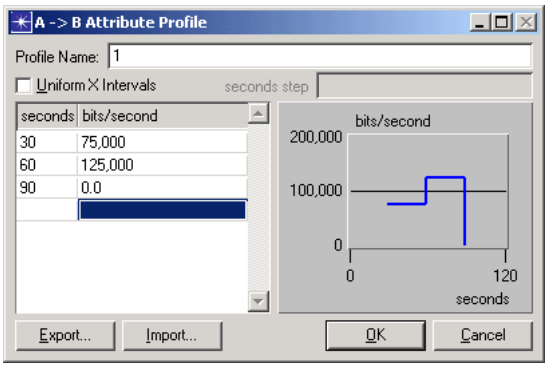

In Figure 8-9 the profile specifies the following (unidirectional) traffic levels on the link:

- 0 to 30 seconds: 0 bits/second
- 30 to 60 seconds: 75,000 bits/second
- 60 to 90 seconds: 125,000 bits/second
- 90 seconds to end of simulation: 0 bits/second
- 120 seconds to end of simulation: 0%

See Traffic Profiles on page ISU-8-5 for more information about profile attributes.

2.2 Configure any additional attributes that define the unidirectional load.

To specify traffic in the opposite direction, repeat steps step 2.1 through step .

End of Procedure 8-6

To create device loads manually, use the following procedure.

Procedure 8-7 Creating Device Loads Manually

1 Right-click on a device (for example, a workstation) and choose Edit Attributes.

➥ The Attributes dialog box opens.

2 Click on the Value field of the attribute that specifies background utilization for that device; then choose Enter from the pull-down menu. The exact name of this attribute depends on the type of object. For example, a LAN node might have a "LAN Background Utilization" attribute, and a server might have a "CPU Background Utilization" attribute.

➥ The background utilization attribute table appears with the default background values.

3 In the Rows field, enter the number of utilization periods you want to use.

➥ The specified number of rows appears in the table.

4 In each row, specify the start time of each utilization period in seconds, the utilization percentage, and the direction in which the utilization applies.

Figure 8-10 Node Background Utilization (Example)

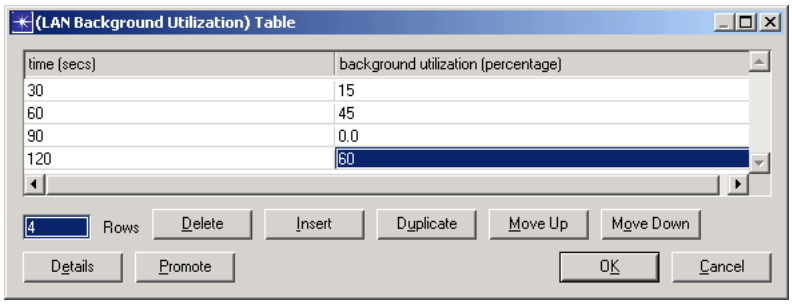

In Figure 8-10 the table specifies the following utilization levels on the LAN:

- 0 to 30 seconds: 0%
- 30 to 60 seconds: 15%
- 60 to 90 seconds: 45%
- 90 to 120 seconds: 60%

- 120 seconds to end of simulation: 0%

End of Procedure 8-7

Working with Background Traffic

Background traffic is incorporated into the standard models in two ways: traffic flows (background routed traffic) and device/link loads (background loads). Each of these two approaches has appropriate applications in a simulation study, and in some cases, they can be used together.

Using Device/Link Loads

Device/link loads rely on the user to specify background loads or utilizations on specific network objects. For example, you might want to specify additional traffic loads on one or more links. As explained previously, this means that a mathematical model is used to estimate the effect of the background loads on these links. These background loads affect the explicit traffic that traverses the links, and the effect is to delay that traffic by emulating additional packets queuing for access to each link. For information about how to specify device and link loads, see Device/Link Loads on page ISU-8-16

Background loads are provided not only for links, but for some other objects as well. These include LAN objects and servers. In each case, the meaning of a device/link load is the same. The resource of interest is receiving an additional amount of usage, either as a traffic intensity (for link objects) or as a percentage of the maximum throughput (for node objects). Keep in mind that device/link loads are considered in addition to other types of traffic, such as traffic flows and explicit traffic. Therefore, a link with a background load of 2 megabits/second and an explicit load is 1 megabit/second during a specific interval has a total traffic load of 3 megabits/second.

Background loads are useful in the following situations:

- You already know the utilization of a component that provides a service to some of your explicit traffic—You want to model this “other utilization” analytically to shorten simulation run times. For example, client traffic traverses two LANs and two routers to reach a server. Server traffic also traverses this path in the reverse direction. The client’s LAN is known to be 80% utilized by a large group of workstations. You may use a LAN object and set its “Background Utilization” attribute to 80% beginning at time zero of the simulation. At the same time, you can set up a client and a server and a router to explicitly model the point of view of a particular user, as depicted below:

Figure 8-11 Using LANs and Background Utilization to Simplify a Network

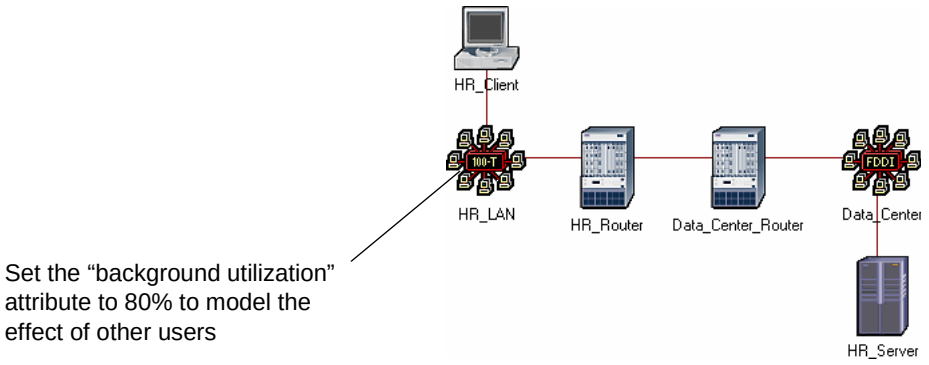

- You don’t know the utilization of a component, but want to determine the effect of increasing it—For example, what will happen to application response time at various points in your network, if you increase load on the server to 60%, from its current value of 30%. In other words, how sensitive is your application to the server’s performance. If it’s not, and network and protocol effects dominate, it may be pointless to upgrade your server. The “background utilization” attribute on a server represents other tasks contending for the server’s processing bandwidth. As this number increases, tasks submitted by your “stakeholder user” are processed less quickly, reflecting the sharing of the server’s resources with others.

Using Traffic Flows

Traffic flows provide another way to specify that the network and its components are processing traffic without having to represent in detail the patterns of that traffic. This approach does not require you to know the amount of traffic incident upon a given component; instead it calculates this amount for you based on an end-to-end description of the traffic. “End-to-end” means that traffic is specified from its originating host, to its final destination.

Traffic flows are specified and generated at the network layer. More precisely, as of this printing, this type of traffic operates only above the IP protocol. That is, every device containing an IP layer is also capable of being a source or destination for end-to-end traffic. This includes routers, servers, and workstations, but also LANs. Switches, for example, are not capable of acting as sources and destinations of end-to-end traffic.

Table 8-2 Sources and Destinations of Traffic Flows

Routers
Workstations
Servers
LANs
End of Table 8-2

Specification of traffic flows essentially consist of a “traffic matrix”, where every host is considered a possible source of traffic for every other host. Here the term “host” is used generically to represent one of the sources or destinations of traffic, as shown in the table above. As you can imagine, this matrix can grow quite large. With just 1000 such hosts, for example, a traffic matrix with one million possible pairs of traffic would be defined. Fortunately, the traffic flow system is intelligent, allocating resources only for source-destination flows that actually contain traffic.

After you configure flows with the specified traffic, the models automatically route this traffic to the relevant devices in the network. This is done by using the routing models, such as RIP, OSPF, and IGRP, the same way explicit traffic is routed. Routing the traffic flows has the effect of creating traffic throughout the network on links, routers and switches, wherever the specified traffic would normally flow. With end-to-end traffic specification (as opposed to component-by-component), changes in the network infrastructure, or even changes like moving a server, are reflected correctly and automatically. This happens because the routes are recalculated in a new simulation of the network. For example, if you specify that 1000 packets per second are to flow between A and B over a defined interval, you can keep this traffic specification for two different network scenarios with a different path connecting A and B. In each simulation, the correct path will be selected and the traffic flows will affect corresponding parts of the network.

Specifying Traffic Flows

Using the traffic flow system brings up an important issue: how can you get this traffic information into the flow objects? Of course, manual entry of traffic profiles between sources and destinations is supported.

Scaling Traffic—Analyzing the Impact of Traffic Growth

A common approach to conducting traffic sensitivity studies is to create end-to-end traffic and scale it to progressively higher values in a series of scenarios. In each scenario, the effect on a particular stakeholder who runs an explicitly modeled application, is observed. The approach of simulating the "stakeholder's" viewpoint is illustrated below:

Figure 8-12 The Viewpoint of a Stakeholder Contending with Background Traffic

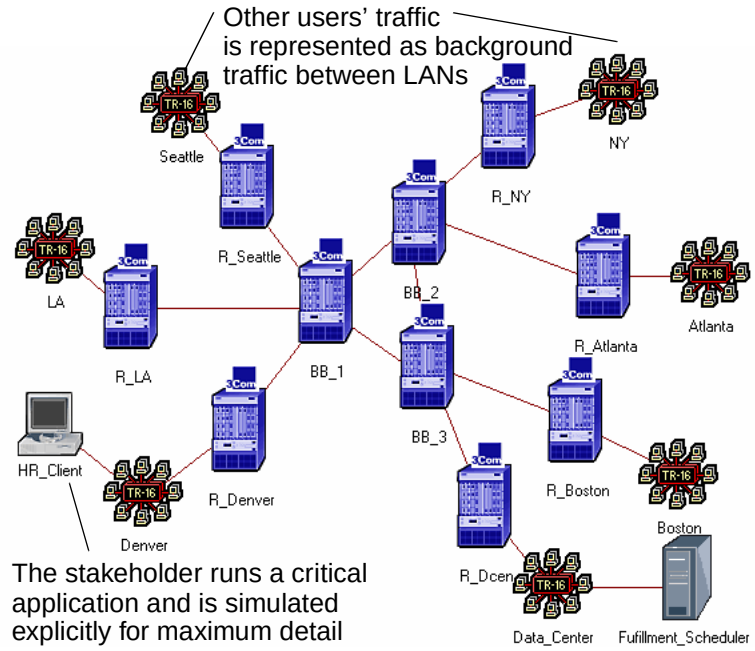

A common approach for a scaled traffic flow study is outlined below:

1) Develop a "baseline" scenario within a new project. You can do this either by importing topology, or by manually constructing the necessary portions of your network's infrastructure. The baseline scenario typically represents the network's normal or present-day infrastructure and traffic.

2) Import traffic for the baseline scenario. Browse the traffic to make sure you understand its nature and that it corresponds to your expectations. Revise the traffic matrix if you feel this is necessary to provide a valid baseline.

3) Specify "stakeholder" traffic for applications of interest. These are one or more end-users of the network who run specific applications. Characterize the end-user's traffic as well as you can based on knowledge of the applications being performed. You may also choose to represent the end-user using measured information, as explained in the SMARTE methodology later in this chapter. In any case, you don't need to create a large number of users, in general; just position users in parts of the network that are of most critical interest, or where you suspect application performance might be problematic.

4) Choose results to be collected. Run the baseline scenario's simulation, and view the results.

5) Create duplicate scenarios and scale traffic. For each scenario, edit the traffic matrix and increase the traffic by an appropriate percentage. When scaling traffic, consider if you want to scale both traffic flows and device/link loads (considered generally to be background traffic) and explicit traffic, or only background traffic. In most cases, scaling just the background traffic is desirable, because the traffic increase is due to more users appearing in the network, not to increased activity by individual users. The purpose of having explicit traffic in the simulation is to represent one or more typical users accurately. Their particular traffic should be negligible when compared to the background traffic. Therefore, not scaling their contribution should not significantly dilute the selected scaling factor.

   Scaling traffic can be performed via the "Traffic Scaling Factor" Simulation attribute. (Recall that simulation attributes are specified under the "Configure Simulation" dialog box in the Project Editor). The "Traffic Scaling Mode" simulation attribute determines if you will scale only the background traffic, or all traffic (including explicitly modeled applications).

6) Run simulations for the duplicated scenarios. Use the "Compare Results" operation to understand how scaling traffic impacts application performance.

Note that if you have purchased the Expert Service Prediction (ESP) option, you may prefer to use the "Specify Traffic Growth" operation in order to quickly run a series of scenarios and collate and compare results in a more automated fashion.

Creating Cross Traffic for a Specific Network Element

Certain network components do not provide a background-load attribute, as do LANs or servers. In particular, as of this printing, router and switch models do not allow you to model the effects of additional load based on an attribute belonging to the device itself. However, representing such a load is sometimes necessary in order to model the "rest of the network" in terms of its effect on the component of interest. If the load submitted to this device is known (for example, from a measurement), it is typically preferable to avoid constructing a full model of all other parts of the network.

If the device of interest has no device/link load attribute, you can still easily specify the desired volume of traffic flowing through the device of interest.

Procedure 8-8 Specifying Traffic through a Device That Does Not Have a "Background Utilization" Attribute

1 "Sandwich" the device between two routers. In other words, create two routers with the correct layer 2 technologies to connect to the device of interest. For instance, if the device is a router with available Token Ring ports, choose two routers that also have Token Ring ports (you can even use the same router model as the device of interest in this case). Connect each router to the device of interest using a bidirectional point-to-point link. These are "artificial routers", used only to generate the background load you seek to introduce for the device of interest. They represent the "rest of the network" from this device's point of view. You can also use LANs instead of routers to represent proxy traffic sources.

Figure 8-13 Proxy Objects Provide Background Traffic Across a Router

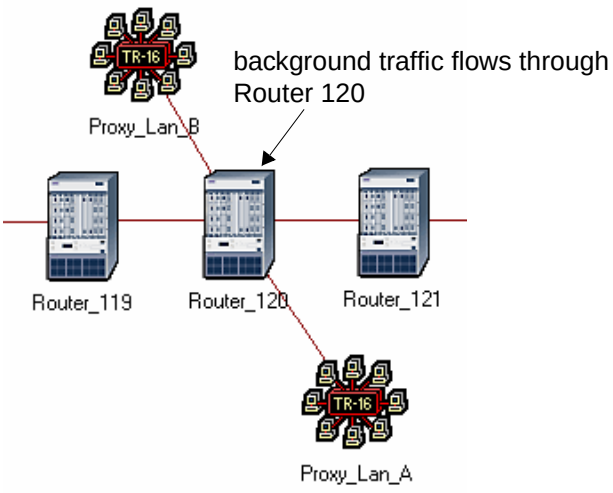

2 Determine the background traffic you want to impose on the device of interest. Devices like routers and switches are mostly affected by the number of packets they process, not the number of bits. Therefore, the volume in terms of packets per second is really what matters. This is a number you will generally obtain from measurement or monitoring tools available in your environment. It may also be an estimate you have made based on statistics collected from other simulations.

3 Specify the background traffic between routers. Which router you choose as the source or destination is not important because your goal is only to make packets flow through the device of interest. Go to the traffic browser to specify the desired volume of packets per second flowing through the device. The number of bits per packet is not as important. Choose a number that represents a typical frame size for the LAN technology being used to connect the devices.

End of Procedure 8-8

There is one caveat to the previous procedure when the device of interest is a router. The issue stems from the fact that certain router models account for concurrent processing of packets on distinct slots. This is controlled by the "IP Processing Scheme" attribute of the router, which is typically available on advanced models. You will need to know if the device under test is a router

which uses a processing scheme which is “Slot Based Processing”. If so, traffic flowing through the router on certain ports may not affect traffic flowing through on other ports. This depends on whether the ports belong to the same slot of the router. Accordingly, if the background traffic setup between the two “artificial” routers is not on the same slot as other traffic of interest, the desired effect may not be obtained. Therefore, in order to accomplish sharing of router resources by the background and the explicit traffic, you may have to carefully control the ports that are selected for each of the links attached to the router. Information on how to do this is provided in the “Working with IP” section of this chapter.

A Collection Strategy For Background Traffic

Common sources for background traffic are RMON 2 probes. Using the Multi-Vendor Import product option, you can also import background traffic from Network Associates’ Sniffer data if it is saved in ExpertSniffer format. Irrespective of the tools you use to collect your information, you will face two other important questions:

- Where to Collect?
- When to Collect?

Positioning your Collection Devices

Let us first address the issue of where to collect. First of all, we should ask ourselves where are all the possible places to collect traffic? The assumption made throughout this section, is that the finest grain location we would consider is the LAN segment. Because of the broadcast nature of local area network traffic on one segment, a collection device is able to observe all traffic sent to or from that segment. Placing more than one collection device on the same segment would result in redundant information.

Note that a LAN segment is not the same thing as a LAN. A LAN segment is a layer-2 physical broadcast domain, such as a shared Ethernet segment, or a token ring, or FDDI segment. If these or other layer 2 technologies are implemented in a switched environment, then each port of a switch corresponds to a distinct segment, because traffic sent on one port of the switch is not observed by stations accessible via the other ports of the switch. In VLAN environments, broadcast domains can be defined logically, but these do not qualify as segments because the switches that implement the VLANs selectively transfer traffic to their different ports.

Our objective is to build a “traffic matrix” to summarize all traffic on the network of interest over a period of time. Therefore, we need to ensure that any relevant traffic is observed by at least one collection device; otherwise it will not be accounted for. Doing this as efficiently as possible requires some knowledge of the traffic patterns. For example, in the scenario below, we may have prior knowledge that A and B send traffic to and from one another but send relatively

insignificant traffic to C and D. And assume a similar statement could be made about traffic between C and D being exclusive with respect to A and B. Then, we could achieve collection of all inter-segment traffic with just two collection devices, as shown.

Figure 8-14 Minimizing the Number of Collection Devices

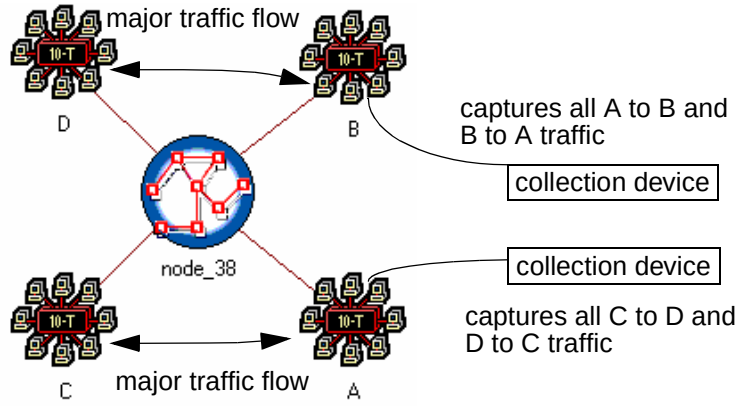

We can achieve this goal because all traffic from A to B must cross segment B and similarly for C with respect to D. In other words, in evaluating traffic volume between two segments, it does not matter on which of the two segments the collection device is positioned: in either case, the traffic will flow through it. So, which traffic flows are not measured with the two collection devices placed as shown? Any traffic between A and C is not measured, because neither has any collection devices attached. If this traffic is truly known to be insignificant for the purposes of our simulation studies, then this traffic collection strategy is acceptable.

The example above serves to illustrate another important point: you must make a decision about whether you want to capture flows that are internal to segments. In the example above, neither A nor B's internal traffic will be reflected in our traffic matrix. If you are primarily concerned about backbone (that is, inter-LAN) traffic, then you can safely follow this approach. You may also be selective in collecting intra-LAN traffic for certain segments, but not others. It is important is that you are aware of the consequences of not placing a collection device at a given segment.

Fortunately, in many networks, we can have the type of prior knowledge we need to efficiently place collection devices. Specifically, highly centralized environments provide this type of advantage. In highly centralized environments, most traffic flows between clients and servers. These servers are concentrated in a relatively small number of facilities. Intense local traffic may

also be present, but that traffic is confined to the local area network (for example, for sending jobs to a local printer). The diagram below illustrates such a network and an effective data collection strategy for characterizing backbone traffic.

Figure 8-15 Collection in a Primarily Centralized Environment

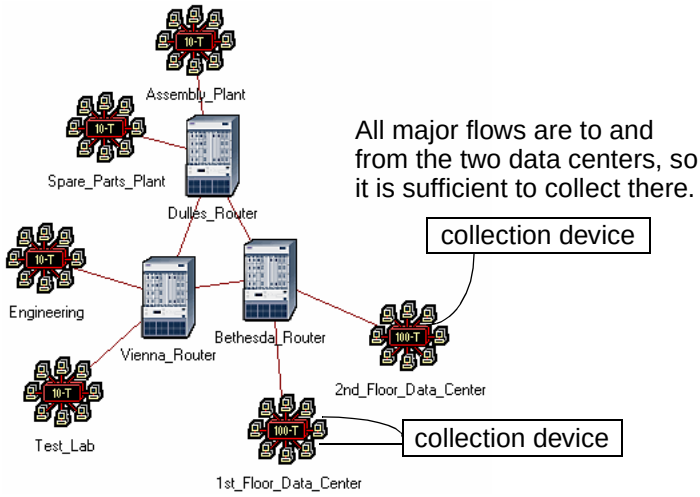

If you are unable to leverage prior knowledge of major traffic flows, as in the examples above, then you may be faced with a more challenging data collection effort, requiring you to do data collection at each segment. However, for characterizing backbone traffic, you may still be able to avoid such widespread collection by looking for aggregation points. Typically, these will be at routers or switches where traffic funnels into the backbone. The diagram below illustrates this approach.

Figure 8-16 Collection in a Decentralized Environment

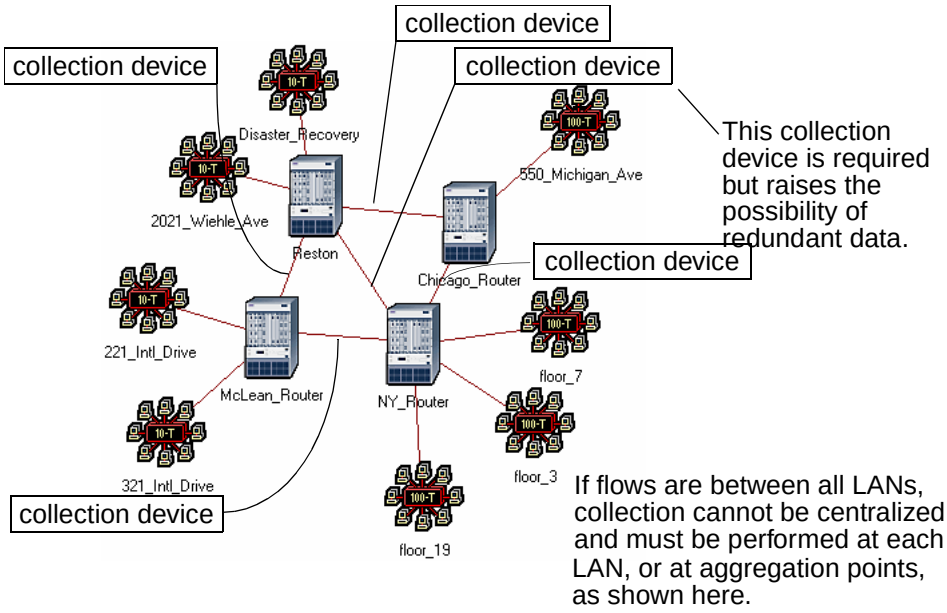

Synchronizing Traffic Collection

One fundamental problem affecting collection is that most organizations possess a limited number of collection devices that they can place throughout the network. Ideally, one of the approaches discussed above permits performing collection concurrently with the number of devices you already have available. However, if this is not the case, you do not necessarily have to acquire new devices in order to collect background traffic.

Your alternative is to collect traffic in phases. For example, if you determine that you require traffic collection at twenty locations in your network, and you own twelve RMON-2 probes. Then, you will need two phases of traffic collection. No individual location should have its traffic collection repeated, so in this case you would utilize the output from twelve probes on the first phase, and from eight probes on the second phase, or some similar partitioning. Relocating probes is not a trivial task, so you probably will not be able to do both phases one right after the other. In this case, you could do the second phase on the following day, but beginning at the same time of day. The objective is to gather data while the network is experiencing similar traffic patterns.

The background traffic system automatically synchronizes all imported traffic information to the same time. The first “bucket” of each source of background traffic is applied at the same time, and that time is determine by the “Background Traffic Start Delay” simulation attribute (set this value in the “Configure Simulation” dialog box).

Managing Redundant Traffic Data

Redundant traffic information typically occurs when collecting in multiple phases. By “redundant”, we mean that two or more distinct traffic profiles are collected for the same pair of hosts, but at different times. If this happens for a particular pair of hosts, it is likely to happen for many in your network.

The traffic collection strategies described above favor collection at the “edges” of the network. The edges are the LANs, or the routers associated with those LANs, if you collect aggregated traffic for several LANs at a time. Collecting at the edges helps minimize redundant traffic information, when compared with collecting in the core of the network (that is, within a router mesh connecting the LANs). However, collection in multiple phases, can cause a segment to have its traffic included more than one phase. Indeed, this is almost impossible to avoid if each segment is being probed, or even with aggregation, because a segment will be directly monitored on one collection phase, but appear as a remote source/destination of traffic, when a data collection device is applied to other parts of the network in a later phase.

The “Import Traffic” feature lets you specify how you want to resolve redundant traffic. You are provided with several options: a) use the higher of the collected traffic volumes at any point over the collection interval; b) use the sum of the traffic; and c) use the lowest recorded traffic level. In practice, option (c) is rarely used, but both options (a), and (b) have appropriate applications. Suppose, for

example, that we have collected traffic twice for a particular pair of hosts, X and Y. Using the higher of the two traffic volumes reflects a conservative, or "worst-case" approach, where the network is assumed to be carrying more load. Approach (b) would be applied if the two traffic flows were considered possible to have occurred simultaneously. This would occur if collection devices were placed within the core of the network and traffic between two hosts can split on to multiple paths, which is possible with certain routing protocols, such as OSPF, or IGRP.

Background Traffic and Statistics

"Background traffic" is considered to be composed of both traffic flows and device/link loads, has one primary purpose: to model the effect of general traffic in the network on selected traffic of interest. This effect primarily takes the form of delay. In other words, background traffic causes additional delay for explicitly modeled traffic. Of course, background traffic causes many other statistics besides delays and response times to be affected. Utilization and throughput on at least some links, for example, increase with additional traffic in the network. Average buffer sizes in routers would also be expected to increase. However, the background traffic system does not cause all of these statistics to be affected. Some statistics do not react to background traffic because of the ambiguities associated with modeling such effects. Others do not, simply because the background traffic architecture is a recent and ground-breaking invention that can still benefit from further integration into the model library. While most important statistics that you will be interested in are integrated with background traffic, the purpose of this section is to alert you to those that are not so that you can comprehensively interpret simulation results. Also, while the statistics enumerated below do not reflect the effect of background traffic, the behavior of the simulation still does in most cases, in terms of injecting delay for explicitly simulated packets. In other words, this is primarily an issue concerning reporting of a specific subset of results, not simulation behavior. See section Limitations of the Background Traffic System on page ISU-8-31 to determine which components do not provide additional delay modeling in response to background traffic.

Table 8-3 Statistics that Don't Reflect Device/Link Load System[1]

Queue and buffer sizes
Global statistics of all types
Server Task Load
Shared Ethernet
Switched and shared token ring and FDDI
LAN Inbound, Local, Outbound, and Through traffic
Statistics of Protocols at higher layers than IP (for example, TCP, UDP, Applications)
End of Table 8-3

1.check for latest information

The fact that certain statistics are not integrated with the background traffic system does not mean that these statistics are not accurate or meaningful. It simply means that these statistics have a different interpretation. Specifically, you should interpret these statistics as reflecting only the activities caused by explicitly simulated traffic (for example, traffic generated by application models specified on workstations and LANs). Finally, as integration of the background traffic system continues, you should check for availability of model library upgrades that would change the list above.

Limitations of the Background Traffic System

As a user of the background traffic system, it is important for you to be aware of its limitations. Certain limitations result from the fact that analytical modeling, and therefore hybrid analytical/discrete-event simulation, is not capable of capturing all of the same behavior as pure discrete-event simulation. Other limitations will be overcome with time as enhanced versions of the background traffic system are developed. Despite these existing limitations, which are summarized in this section, the background traffic system is extremely useful, providing an approach to reduce simulation run times while still representing important effects of traffic in the network.

- The background traffic system is IP-centric in its current implementation. This means that devices that are the sources and destinations of background traffic must contain the IP protocol. Pure IPX systems, for example, are not able to act in this capacity. Because IP is essentially the “conduit’ for background traffic information, only intermediate devices that support IP or that support IP at lower layers can be affected. Thus a Frame Relay switch that is part of an overall IP infrastructure can carry background traffic. However, no background traffic architecture exists at this time for a pure Frame Relay network that does not serve as a bearer of IP traffic.

- Traffic flows (background routed traffic) only impact the network, not the server. Generating extra traffic in the network, by scaling the traffic flows for example, will not create a corresponding increase in task processing time on the server. This is due to the fact that the background traffic system is not capable of determining how traffic correlates to processing load on servers. It is therefore up to the user to take steps to properly load the servers. One way to accomplish this is to change the servers’ “background utilization” attribute.

- Background traffic is transparent to certain protocol mechanisms. While many protocols in the model library are aware of the background traffic system, not all of their functionality reflects the presence of the additional traffic. For example, the weighted fair queuing mechanism of IP does not estimate background traffic’s contribution to queue lengths within each priority, because background traffic does not carry priority attributes at this time. Additional protocols and mechanisms that are not sensitive to the “background utilization” attribute are listed in the table below. This information is valid as of this printing. As integration of the background traffic system continues, you should check for availability of relevant model library upgrades.

Table 8-4 Models/Mechanisms that Do Not Account for Traffic Flows[1]

Non-LAN objects using switched or shared Token Ring
Non-LAN Objects using switched or shared FDDI
Non-LAN objects using half-duplex (shared) ethernet
End of Table 8-4

1.check for latest information

9 Customizing Models

Subnetwork, node, and link objects in IT Sentinel have *attributes* that define the object and control its behavior in a network model. Although the set of attributes associated with an object can vary, they all define the behavior of an object or model.

Attributes

IT Sentinel allows you to manipulate attributes in many ways, including:

- setting attribute values for one object or a group of objects
- defining attribute properties
- saving attribute properties for reuse with other attributes

Three types of attributes may comprise an object. These types are built-in attributes, model attributes, and extended attributes.

- Built-in attributes provide basic information about an object. They specify information such as the object's name, location in the workspace, and graphic representation. Built-in attributes are characteristic of each object type. When you create an object, IT Sentinel automatically specifies initial values for these attributes.
- Model attributes add information to a node or link object. They determine how the object behaves in a network model by specifying protocols, data rates, and similar information. Model attributes automatically become part of any object using that model. The model developer defines model attributes and sets initial values for them.

All attributes share a set of basic *properties*. Attribute properties define the attribute, specifying information such as its name, data type, and allowed values. These properties are generally fixed, but can be given different values in some cases. For information on attribute properties, see Attribute Properties on page ISU-9-18.

Attributes Dialog Box

You view and set attribute values through an object's Attributes dialog box. For each attribute of an object, this dialog box displays the name, value, and units (if any) within a data table.

Every Edit Attributes dialog box has essentially the same structure, regardless of the type of object to which it applies. When you choose Edit Attributes from an object's pop-up menu, the dialog box displays all *primary attributes* associated with the object. Primary attributes are those you use most often to define an object. For example, primary attributes specify the object's name and its underlying (node or link) model.

Network objects (subnets, nodes, and links) also have *advanced attributes*, which specify parts of the object definition that you might set less frequently. Advanced attributes typically define information such as an object's appearance and location in the workspace. To view the advanced attributes of an object, choose Advanced Edit Attributes from the object pop-up menu or set the Advanced check box in the Edit Attributes dialog box.

Note—Primary and advanced attributes exist only for network objects. For all other OPNET objects, this dialog box displays all of the object's attributes.

Figure 9-1 Attributes Dialog Box

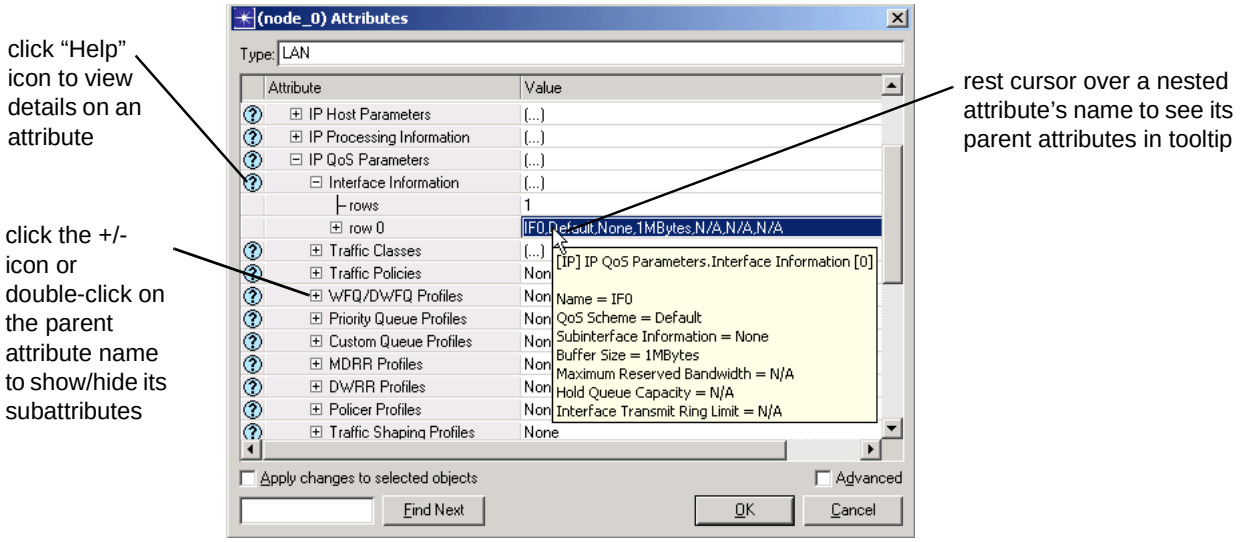

The Edit Attributes dialog box appears whenever you right-click on an object and choose Edit Attributes or Advanced Edit Attributes. This dialog box has the same general interface for nodes, links and configuration objects.

- To expand or collapse a compound or grouped attribute, double-click in the Attribute field.
- To edit or view a compound attribute in a new window:
  - — If the Value field contains an ellipsis ((...)), double-click in the Value field to edit it.
  - — Otherwise, click in the Value field and choose Edit... from the menu.
- To view the parents of a nested attribute, rest the cursor over the attribute's name (that is, in the Attribute column); the full attribute name appears in a tooltip.
- To view details on an attribute, click on the icon in the leftmost column.
- To search for a particular attribute, enter a string in the Find field (lower left corner) and click Find Next. IT Sentinel finds the next instance of an attribute whose name or value contains the search string.

In addition to attribute names and values, the Edit Attributes dialog box also provides buttons for various operations such as viewing attribute properties. These operations vary depending on the type of object the dialog box supports. The following table describes the operations available in the Edit Attributes dialog box.

Table 9-1 Attributes Dialog Box Operations

Operation	Access	Description
Apply Changes to Selected Objects	Check box	Applies changes to an object's attribute values to both the object and all selected objects in the workspace. The rules for applying these changes are discussed under Setting Attribute Values on page ISU-9-4.
Find	Edit field/button	Search for attributes with names or values that contain the specified search string
Extended attrs.	Button	Summons the Extended Attributes dialog box, which allows you to create new attributes for the object.
Link Interfaces	Button	Displays interface and port information for nodes connected to the selected link.
Packet formats (selected links only)	Button	Allows you to specify which packet formats a link supports. (This button is available only for links that permit you to set packet formats.)
Promote	Button	Promotes the attribute so that its value can be specified at a higher level in the network model or when the simulation is executed.
Redefine path (links only)	Button	Allows you to redefine a link's path by redrawing it in the workspace.
Advanced (subnets, nodes, and links only)	Check box	When selected, displays all of the attributes (primary and advanced) associated with the object. By default, only primary attributes are displayed.
Details	Button	Summons the Attribute Properties dialog box, which displays detailed information about the selected attribute.
End of Table 9-1		

IT Sentinel includes three preferences (Edit > Preferences) you can set to highlight attributes that have changed and do not reflect the default settings in the underlying model. These preferences are:

- "mark_nondefault_attrs"—Enable/disable attribute highlighting.

- ”mark_nondefault_attrs.intended_color”—This color indicates that you recently changed an attribute during the current session of editing attributes (that is, after you opened the Edit Attributes dialog box).
- ”mark_nondefault_attrs.changed_color”—This color indicates an attribute that was changed (by a user, the program, or a simulation) before you opened the Edit Attributes dialog box. As a result, the attribute might not reflect the default value specified in the underlying model.

In addition to attribute names and values, the Attributes dialog box also provides buttons for various operations such as viewing attribute properties. These operations vary depending on the type of object the dialog box supports. The following table describes the operations that may be available in the Attributes dialog box.

Setting Distributions

Some attributes take a distribution as a value. Distributions are used to generate a series of random numbers with specified characteristics (such as normal or constant). When you edit the value of a attribute that uses distributions, a dialog box opens that allows you to select the type of distribution to use and enter values for one or two parameters that define the shape of the distribution (see Figure 9-2 for an example). The predefined distributions supplied with IT Sentinel are described in Chapter 6 Distributions on page ISR-6-1.

Figure 9-2 Distribution Specification Dialog Box

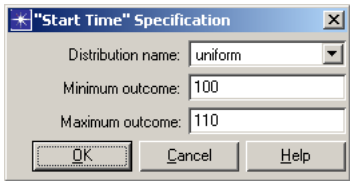

Setting Attribute Values

You can set attribute values for one object or for several objects at once. You can even apply changes to different types of objects; however, they must contain the same attributes. For example, subnetwork and node objects both contain the “y position” attribute. Changes made to this attribute in one object can be applied to the other (to align them in the workspace, for example).

You set an object’s attribute values through its Attributes dialog box. Depending on the attribute, you might set its value in any of several ways. For example, you might type a value into a text entry area, select a value from a pull-down menu, summon a secondary dialog box, or select from a palette. Methods of setting attribute values are identified with the corresponding attribute discussions throughout this manual.

Setting Attributes for One Object

To set the value of an attribute on a single object, use the following procedure.

Procedure 9-1 Setting Attribute Values for One Object

1 Move the cursor over the desired object and choose Edit Attributes or Advanced Edit Attributes from the Object pop-up menu.

- Edit Attributes opens the object's Attributes dialog box with the primary attributes visible.
- Advanced Edit Attributes opens the dialog box with all of the object's attributes (both primary and secondary) visible.

From either version of the dialog box, you can change the list of attributes displayed by selecting or deselecting the Advanced check box.

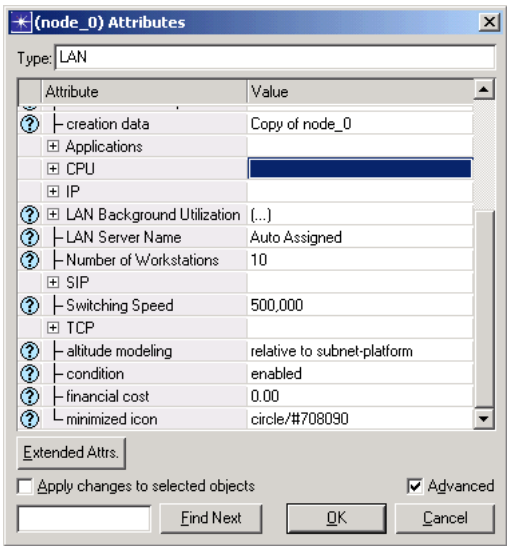

2 Set an attribute value by clicking in the appropriate cell in the Value column. The editing method may require text editing or selection from a pull-down menu.

3 Click OK to save the changes and close the dialog box.

End of Procedure 9-1

Setting Attributes for Multiple Objects

You might need to set values for the same attributes on several objects in a network model. For a few attributes, you could make these changes individually using Procedure 9-1. However, this method can be tedious and error-prone for large numbers of objects. IT Sentinel provides several alternate methods for quickly changing attributes on many objects.

- Using the Apply Changes to Selected Objects check box in the Attributes dialog box.

- Using the Edit Selected Objects operation on the workspace pop-up menu.
- Using the Edit Similar Nodes/Links operation on the object pop-up menu.
- Using the Edit Objects Using Template operation on the Edit menu.

These methods are described in the following sections.

Using the Apply Changes to Selected Objects Method

With this method, you manually select every object to be changed, specify the changes for one object, then direct that the changes be applied to all selected objects. IT Sentinel applies the changes to the corresponding attributes of every selected object.

When applying attribute changes to several objects, IT Sentinel begins by noting every change to the one object whose attributes you are setting. It then tries to make corresponding changes to every selected object, one attribute at a time. In each case, the change is applied if all of the following criteria are met:

- The attributes being set must have the same name and type.
- The new attribute value must agree with the attribute properties of the attribute in the selected object. For example, the new value must be within the range of allowed values for the selected object's attribute.
- The new attribute value must pass all rules governing attribute values, such as the requirement for unique node names within a subnet.

Note—If some of the attribute changes fail, IT Sentinel changes only those attributes that passed the preceding tests. IT Sentinel does not warn you about any failed attribute value changes. If necessary, open the Attributes dialog box of a selected object to verify the changes.

Some additional rules apply when extended attributes are involved:

- Any extended attribute deleted from the object being set is deleted from all selected objects.
- Any extended attribute appearing on the object being set, but not on a selected object, is added to the selected object.

- Any extended attribute appearing on both the object being set and a selected object is retained on both. If the properties of such an extended attribute are changed, the following additional rules apply:
  - — Any changes to the properties of an extended attribute in the object being set are also made to the selected object.
  - — If a property change causes a previously-set value to become invalid, that value is removed and either promoted or replaced with its default value, depending on the Auto. Assign Value setting of the extended attribute.

WARNING—If you change a sub-attribute of a compound attribute, the entire compound attribute value will replace the corresponding compound attributes on the selected objects. This can produce unwanted changes in other sub-attributes of the compound attribute. To edit only specific sub-attributes, use the procedure given in Using the Attribute Template Method on page ISU-9-9.

Procedure 9-2 Setting Attribute Values for Several Objects

1. Select the objects whose attribute values will be set. You can select the objects manually or automatically, as described under Logical Object Selection on page ISU-9-15.

   ➥ Selection markers appear around the objects.

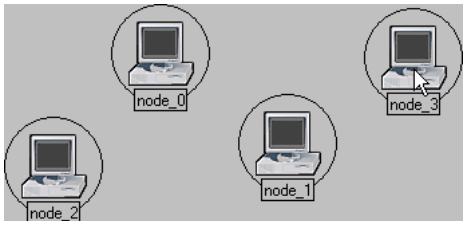

2. Open an object's Attributes dialog box. This object does not have to be selected, because changes always apply to the object whose Attributes dialog box is open.

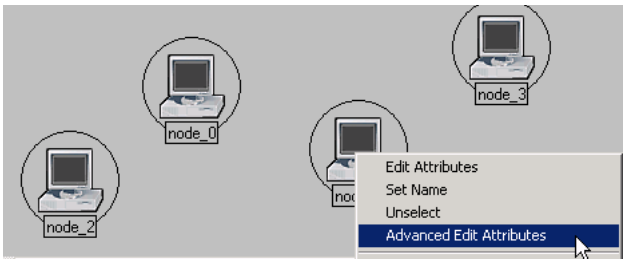

3 Set the desired attribute values.

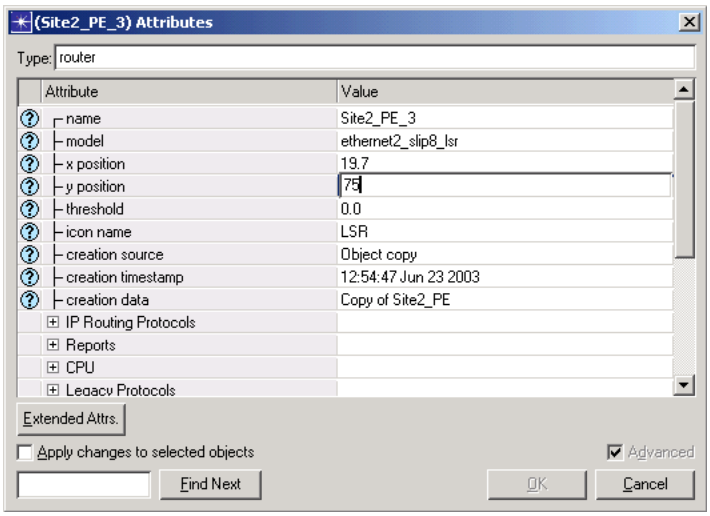

4 Select the Apply Changes to Selected Objects check box, then click OK.

➥ The Attributes dialog box closes and the attributes of the selected objects are updated.

In this example, the objects all move to the same y position.

End of Procedure 9-2

Using the Edit Selected Objects Method

This method provides an alternate way to edit the attributes of several objects at once. IT Sentinel checks to see what attributes the objects to be edited have in common and creates a table of these attributes. You can then edit the attribute values as needed.

Procedure 9-3 Setting Attribute Values on Selected Objects

1 Select the objects whose attributes you want to change. They do not have to be the same type of objects.

2 Right-click in the workspace and choose Edit Selected Objects from the workspace pop-up menu.

➥ The Objects Attributes dialog box opens, displaying all the attributes that the selected objects have in common.

3 Change the desired attribute values in the Objects Attributes dialog box. When you are finished, click OK.

➥ The object attributes are updated appropriately.

End of Procedure 9-3

Using the Edit Similar <Nodes/Links> Method

This method is similar to the Edit Selected Objects method except that it automatically selects the set of objects to consider, based on the value of their "model" attribute.

Procedure 9-4 Editing Attribute Values on Similar Objects

1 Right-click on an object in the workspace, and choose Edit Similar Nodes/Links from the object pop-up menu.

➥ The Attributes dialog box opens, displaying all the attributes of all the objects with the same "model" attribute as the selected object.

2 Change the desired attribute values as appropriate. When you are finished, click OK.

➥ The object attributes are updated.

End of Procedure 9-4

Using the Attribute Template Method

This method generates an editable table of attribute values based on a user-defined attribute template (for an example, see Figure 9-3). Because the table is limited to a set of interesting attributes in the network, for all or only certain objects, you can change the attribute values quickly and easily.

Figure 9-3 Edit Objects by Template Dialog Box

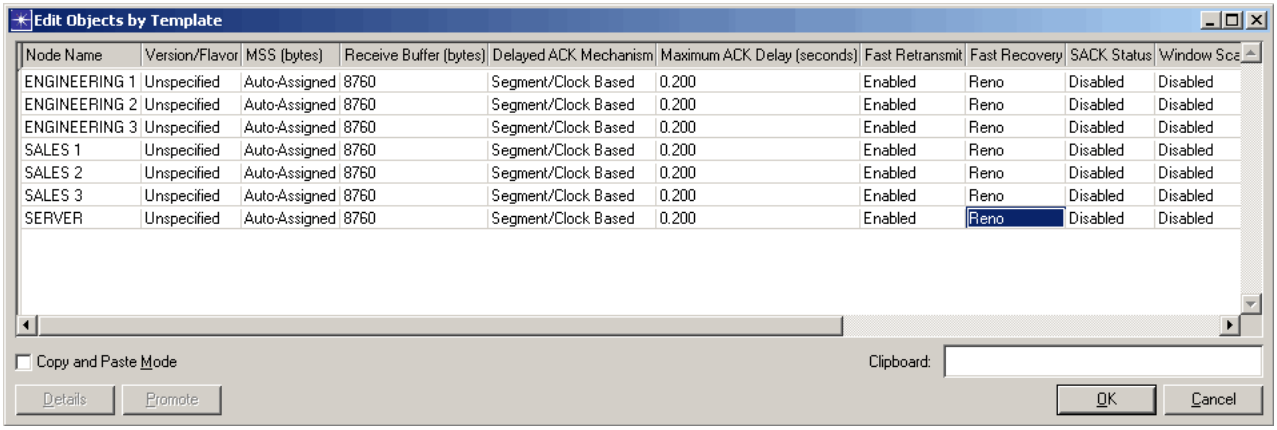

Procedure 9-5 Setting Attribute Values by Template

1. Choose Edit > Edit Objects Using Template.

   ➥ The Edit Attribute Template dialog box appears.

2. Select an attribute template from which to generate an attribute table.

   If you want to define a custom template, click Close, define a template as described in Procedure 9-6, then begin this procedure again.

3. Set the “Include only selected objects in report” and “Ignore views” checkboxes to control which objects will be included in the attribute table. The default settings will include all objects in the network.

4. Click Generate.

   ➥ The Edit Objects by Template dialog box appears.

5. Edit the attributes as needed. There are two editing modes:

   Standard Mode

   You can click on any attribute value and edit it as you would in the Attributes dialog box (see Procedure 9-1 Setting Attribute Values for One Object on page ISU-9-5).

   Copy and Paste Mode

   You can copy an attribute value and paste it to one or more similar attributes. This mode disables the standard methods for changing attribute values (edit fields, pull-down lists, and so on).

   5.1 Select the Copy and Paste Mode checkbox.

   5.2 Select an attribute value and press Ctrl+C.

   5.3 Select the attribute value or a range of contiguous values and press Ctrl+V. (You can paste only into the same attribute column that you copied from.)

6. When you are finished editing attribute values, click OK to make the changes.

End of Procedure 9-5

Defining an Attribute Template

Attribute templates are used to generate tables of specific attribute values for one or more objects in a network. These attribute tables are used when you:

- Edit attribute values for multiple objects (see Using the Attribute Template Method on page ISU-9-9).

- Create custom reports of attribute values for screen or web viewing (see User-Defined Reports on page ISU-11-3).
- View a real-time table of selected attribute values (see User-Defined Reports on page ISU-11-3).

An attribute template specifies the following information about an attribute table:

- Table category and name—Every attribute table belongs to a category, such as “IP” or “Application”, and has a name that identifies it.
- Object types to include—A table can include all network objects, a specific type of network object, or objects selected by custom criteria that you specify.
- Attributes to include—Each row of an attribute template specifies the title and contents of a column that will appear in the attribute table.

To define a new attribute template or edit an existing one, follow Procedure 9-6.

Procedure 9-6 Defining or Editing an Attribute Template

1 Choose Edit > Edit Attribute Template.

➥ The Define Attribute Template dialog box appears. Each table row in this dialog box corresponds to an attribute column in tables generated from this template. At first, the only column defined is for an object name.

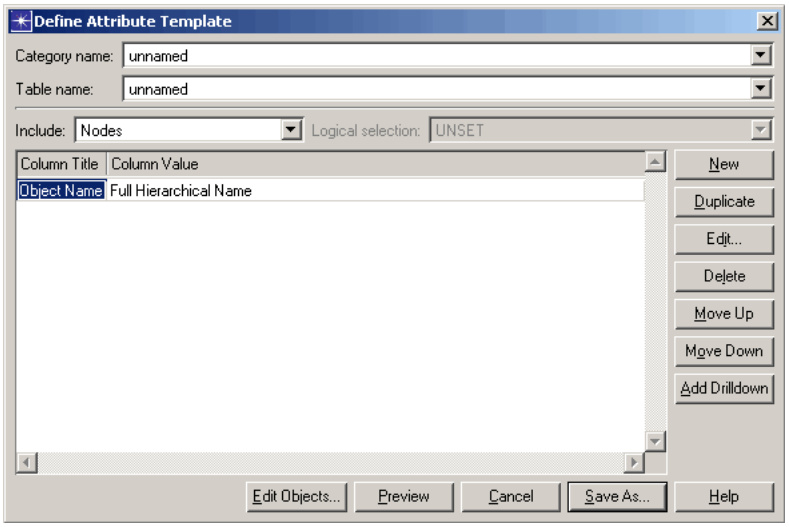

2 Choose whether to define a new template or edit an existing one:

- To define a new template, leave the category and table names set to “unnamed”. (You will specify these names when you save the template.)
- To edit an existing template, select its category and name from the pull-down menus.

3 Choose what object types to include in the table.

3.1 Select an object type from the Include pull-down menu.

3.2 If you selected “Objects in logical selection”, use the Logical selection pull-down menu to specify a selection criterion. You can use one of the predefined criteria or create your own (as described in Logical Object Selection on page ISU-9-15).

Note—If you define a new selection criterion, you must close the Define Attribute Template dialog box and return to Step 1 to make your criterion appear in the Logical selection pull-down menu.

4 Choose what attributes to include in table columns.

- To add a new undefined column, click New.
- To add a new column based on an existing one, select the desired attribute row and click Duplicate.
- To edit a column definition, select a row and click Edit. (Alternatively, you can double-click the Column Value field.) This opens a dialog box for specifying the contents and format of the table column; see Editing an Attribute Table Column on page ISU-9-13 for details.
- To delete a column, select its row and click Delete.
- To add a drilldown, select a row and click Add Drilldown.

  A drilldown is a child table that you open by following a link in the parent table. Drilldowns are useful when you want to show a set of related attributes in a separate window. For example, you might want to create a drilldown table for IP port attributes on a router.

  This operation creates a drilldown table definition, which is initially equivalent to the current table definition. To edit the drilldown table, select the row and click Edit.

You can also move attributes around to modify the appearance of tables generated from the template:

- To change the sequence of attributes, select an attribute row and click Move Up or Move Down (thus moving the resulting table column left or right, respectively).
- To see what the generated table will look like, click Preview.

5 Click Save As.

- If you are defining a new template, enter names for the category and table. The category can be one of the existing categories or a new one.
- If you are editing a template, you can either keep the original category and table names (to replace the original template) or change the table name (to create a new version of the original).

➥ The template is saved with the extension `.urep.xml`. New templates are saved in your default model directory; edited templates are saved in the same directory as the original template.

6 Click Cancel to close the dialog box.

End of Procedure 9-6

Editing an Attribute Table Column

You use the User Report Column Definition dialog box (Figure 9-4) to specify the contents and appearance of a column in an attribute tables. The dialog box lets you pick an attribute from a list of available attributes. You can assign a column title and specify how the attribute value should be shown. Follow Procedure 9-7 to edit the definition of an attribute table column.

Figure 9-4 Define User Report Column Dialog Box

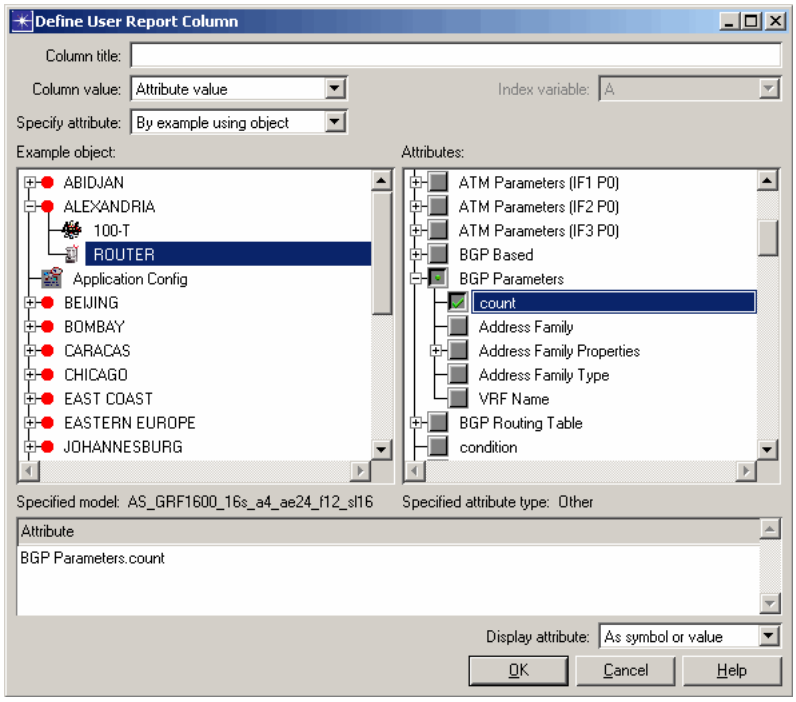

Procedure 9-7 Editing an Attribute Table Column Definition

1 Use the Specify attribute pull-down menu to specify how to view the available an attributes. The options are:

- By example using object—The Example treeview lists all objects in the current scenario.
- By example using model—The Example treeview lists all models in the model directories (specified by the mod_dirs preference).

2 In the Example treeview, select an object or model that contains the attributes you want to view.

➥ The Attributes treeview displays a list of all attributes for the selected object or model.

3 Select an attribute in the Attributes treeview.

4 In the Column title field, enter a title for the table column. (This will generally be some version of the attribute name.)

5 Use the Column value pull-down menu to specify the contents of each table cell in the resulting column. The options are:

- Full hierarchical name—The object name, including all parent objects (for example, `top.subnet_1.node_3`).

Note—Full hierarchical name applies only to an object, so attribute selection is disabled when you select this option.

- Simple name—The object name (for example, `node_3`).

Note—Simple name applies only to an object, so attribute selection is disabled when you select this option.

- Attribute value—The current value of the attribute.
- Parallel attribute values—The current value of one or more similar attributes.

  This option is useful when you want to show related attributes in the same column. For example, a router model might contain compound attributes for three types of interfaces (loopbacks, physical interfaces, and subinterfaces). Each interface attribute includes a sub-attribute that specifies the IP address. Using this option, you could create a column called “IP Address.” You could then select the IP address sub-attributes for the three compound attributes. In the resulting table, the “IP Address” column would show the addresses for loopbacks, physical interfaces, and subinterfaces.

  If you select this option, you should next select some additional attributes in the Attributes treeview.

- Index variable value—The index value of the attribute. This option is useful if a model can contain multiple instances of an attribute. If you select this option, you must also select something from the Index variable pull-down menu.

6 Use the Display attribute pull-down menu to specify how to show the attribute value. The options are:

- As symbol or value—If the model uses a symbol map for the attribute, the symbol is displayed when appropriate; otherwise, the attribute value is displayed.
- As actual value—The attribute value is always displayed.

7 Click OK to accept the column definition and return to the Define Attribute Template dialog box.

End of Procedure 9-7

Selecting Objects by Attribute

In addition to selecting objects graphically (by clicking on them), you can select objects in two other ways:

- Similar object selection—selects objects by model attribute
- Logical object selection—selects objects by type, location, and attributes

These operations can be particularly helpful during group attribute assignment, by simplifying the selection of objects throughout multiple levels of a network model.

Similar Object Selection

This operation allows you to select all nodes or links in a network model that use the same model (that is, have the same ”model” attribute value) as a specified object.

Procedure 9-8 Selecting Similar Objects

1 Choose Select Similar Nodes or Select Similar Links from an object pop-up menu.

➥ The program selects all objects in the network model that have the same “model” attribute value as the object used to open the pop-up menu. Any previously selected objects remain selected.

End of Procedure 9-8

Logical Object Selection

This operation allows sophisticated control over object selection. You can invoke this operation at any time, regardless of whether objects are already selected, and use it to select or deselect objects, based on a variety of selection criteria. Logical object selection has two primary advantages over graphical selection:

- It allows you to select objects that are not visible in the workspace (that is, objects that are nested within subnetworks).
- It allows you to define a set of criteria on which to base the object selection.

Note—If you have objects selected in different subnets, moving between subnets in the Project Editor can clear the object selection set. You might want to use the Network Browser to traverse selected objects in different subnets, because this will retain your selection set. See Network Browser on page ISU-7-18 for information on using the Network Browser

Selection criteria comprise the type of objects to be selected, their locations within the network model (the search scope), and an optional set of attribute values that either may or must be met. You enter these selection criteria via the Logical Object Selection dialog box.

Procedure 9-9 Selecting Objects Logically

1 Choose Select Objects… from the Edit menu.

➥ The Logical Object Selection dialog box appears.

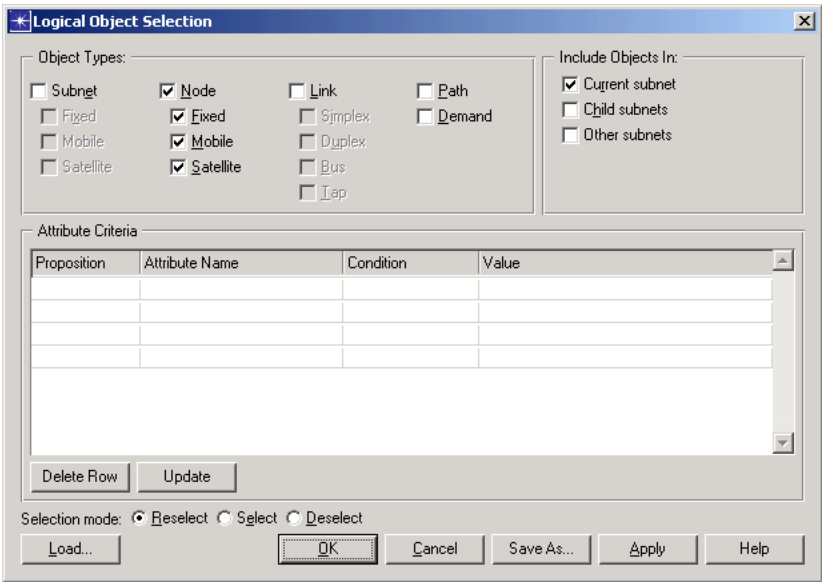

2 Specify the object types to be selected by selecting the appropriate check boxes. You can specify more than one type of object.

3 Specify where in the network model to search for the specified object types. By selecting an appropriate combination of check boxes, you can search anywhere from one subnet to the entire network model.

Table 9-2 Search Scope

Location	Description
Current subnet	Searches only the subnet in the workspace.
Child subnets	Searches all child subnets of the current subnet.
Other subnets	Searches all subnets except the current subnet.
End of Table 9-2	

➥ Based on the object type and search scope, IT Sentinel constructs a list of possible attributes and values.

4 If desired, specify one or more attribute value criteria for the objects to be selected. These criteria are defined in the Attributes data table as follows:

Table 9-3 Attribute Value Criteria

Column	Field Type	Description
Prop. (proposition)	Pull-down menu	Specifies whether to “Consider” or “Require” the attribute value criteria when selecting objects. To be selected, an object must pass both of these tests: • It must meet every attribute value criterion marked “Require”. • It must meet at least one of the attribute value criteria marked “Consider”. (Thus, one “Consider” criterion in the table is equivalent to a “Require” criterion.)
Attr. Name	Pull-down menu	Specifies the name of an attribute that the selected object must contain. This menu lists all visible attributes for the specified object types and search scope. Changing either of these criteria may cause the list of attributes to change.
=?	Pull-down menu	Specifies the required relationship between the attribute and its value: equal to (=), not equal to (!=), greater than (>), less than (<), greater than or equal to (>=), less than or equal to (<=).
Value	Varies, depending on attribute type	Specifies the desired attribute value. Depending on the attribute data type, you can type in a value or select from a list of values.
End of Table 9-3		

5 Specify a selection mode (by left-clicking its radio button) to indicate how IT Sentinel should treat selected objects. Only one selection mode can be selected at a time.

Table 9-4 Selection Modes

Selection Mode	Description
Reselect	Deselects all objects in the specified search scope, then selects only those objects that meet the specified criteria.
Select	Selects all objects that meet the specified criteria. If some objects in the search scope were already selected, they remain selected.
Deselect	Deselects all objects that meet the specified criteria.
End of Table 9-4	

6 Left-click Apply (to select without closing the dialog box) or OK.

➥ IT Sentinel creates a list of all specified object types in the search scope, removes from the list any objects not meeting the attribute value criteria, and selects (or deselects) the remaining objects according to the selection mode. If necessary, the network view changes to show the subnet that contains the first object in the list.

Figure 9-5 Logical Object Selection Example

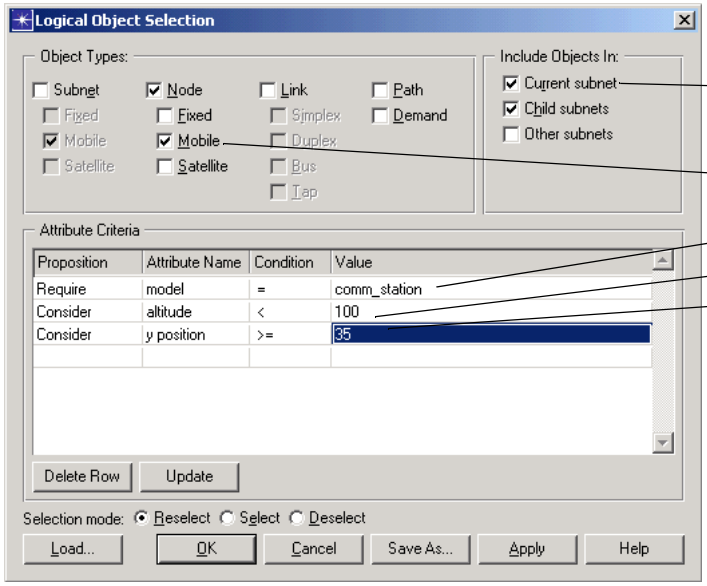

The selection criteria specified here produce the following result:

- In the current subnet and all its child subnets,

all mobile nodes

- using the" comm_station" model AND…
- having an "altitude" less than 100 OR a "y position "greater than or equal to 35…

are selected. All other objects within the search scope are deselected.

Attribute Properties

Every attribute has a set of properties that defines it. The basic properties of an attribute are its "name", "data type", "default value", and "units". Additional properties can provide a description of the attribute, allowed or typical values, and other information. Attribute properties are predefined and usually unchangeable, although you can view them in several dialog boxes. You can change attribute properties at certain times, however, such as when creating an extended attribute.

The following table describes the properties that define an attribute.

Table 9-5 Attribute Properties

Property	Description
Name	Uniquely identifies the attribute during model specification and simulation. Attribute names may not contain a dot "." or left square bracket "["character.
Data type	Indicates the kind of information the attribute can store.
Default value	Specifies a value to be assigned to the attribute when its parent object is created, if the attribute property's Auto. Assign Value check box is selected. The default value may also be used at simulation run time if the attribute has not been assigned another value.
Units	Helps users of a model interpret the attribute's definition, but plays no role during simulation.
Range	Defines a range of allowed values for the attribute. Range applies to numeric attribute values only.
Comments	Describes the attribute and how it should be used.
Properties type	Indicates whether the attribute properties are *public* or *private*. Private properties apply to one attribute only. Public properties can be saved and applied to other attributes. See Sharing Attribute Properties on page ISU-9-20 for details.
End of Table 9-5	

Viewing Attribute Properties

You can view or change attribute properties in the Attribute Properties dialog box. Because you can change properties only at certain times, however, there are two versions of this dialog box. A read-only version restricts you to viewing the attribute properties, whereas the standard version allows both viewing and editing (as described in Attribute Properties Dialog Box on page ISU-9-28). For viewing attribute properties, this dialog box can be accessed several ways:

- Via the Details button in the Attributes dialog box.
- Via the Details button in the Rename/Merge Attributes dialog box.
- Via the Edit Properties button in the Extended Attributes dialog box.
- Via the Edit Properties button in the Derive New Model dialog box.

Procedure 9-10 describes how to access the Attribute Properties dialog box using the Details button.

Procedure 9-10 Viewing Attribute Properties

1 Open an Attributes dialog box (as described under Attributes Dialog Box on page ISU-9-1).

2 Left-click an attribute name to select it.

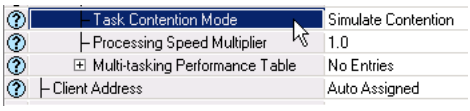

3 Open the Properties dialog box for the attribute:

— If you are in the top-level Attributes dialog box, click the Help icon next to the attribute name.

— If you are in an attribute table, click the Details button. (This button is available only when an attribute is selected.)

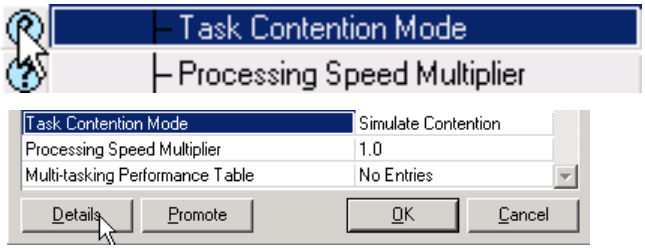

➥ The read-only version of the Attribute Properties dialog box pops up, displaying detailed information about the attribute.

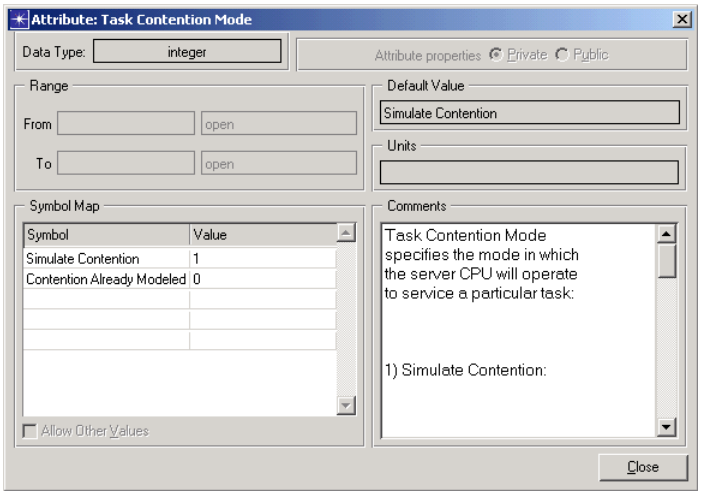

End of Procedure 9-10

Sharing Attribute Properties

Attribute properties are either *private* or *public*, as indicated by the radio buttons at the top of the Attribute Properties dialog box. When private, attribute properties apply only to the attribute for which they are specified. All attributes are private by default.

When attribute properties are public, you can save them to disk. This lets you create a set of properties that can be assigned to other attributes, eliminating the need to respecify the same properties each time they are needed. In addition, because public attribute properties are shared, you can change them on one attribute and all attributes using that set of public properties will automatically inherit the same changes.

Note—Public attribute properties cannot be inherited from a parent model (as described under Inheritance on page ISU-9-29). Thus, if you make attribute properties public, all Inherit check boxes are deselected automatically. They will remain deselected even if you later change the properties back to private.

Attribute properties can be made public or private, saved, and loaded only from the standard (editable) version of the Attribute Properties dialog box. You can switch attribute properties between public and private as needed (for example, to edit public properties).

Procedure 9-11 Saving Public Attribute Properties

1. Select the Public radio button in the Attribute Properties dialog box.

   ➥ The file name **`unnamed`** appears following the Public radio button, indicating that the attribute properties have not yet been saved, and the Save public button becomes active.

2. Left-click the Save public button.

   ➥ A dialog box pops up for naming the set of attribute properties.

3. Type in a name and click OK.

   ➥ The new name appears following the Public radio button. The attribute property set is saved in your primary model directory, appended with the suffix "**`.ad.m`**".

End of Procedure 9-11

Procedure 9-12 Loading Public Attribute Properties

1. Click the "Load public" button in the Attribute Properties dialog box.

   ➥ A menu of available public attribute properties pops up.

2. Select a set of public properties from the list and click OK.

   ➥ The attribute properties in the selected set replace those in the Attribute Properties dialog box and the dialog box becomes view only.

End of Procedure 9-12

Editing Public Attribute Properties

Only private attribute properties can be edited. Therefore, you must switch public properties to private before changing them. You can then switch them back to public and resave them.

Procedure 9-13 Editing Public Attribute Properties

1. Select the Private radio button.

   ➥ The fields in the Attribute Properties dialog box become editable. The Public check box and the Save Public button become inactive.

2. Edit the attribute properties as desired.

3. Select the Public radio button.

   ➥ The original file name appears following the Public radio button, as a default value, and the Save public button becomes active.

4. Save the changed attribute properties, using the same or a new name. If you use the same name, all attributes using this set of properties will receive the changes.

End of Procedure 9-13

Compound Attributes

A *compound attribute* is a grouping of several attributes into one attribute. This grouping allows you to describe a complex set of characteristics for an object with one value. You can do this by defining the compound attribute, then editing its properties to create a symbol map in which each set of characteristics is represented by one symbolic name. Each attribute of a compound attribute has a value and properties that can be viewed or edited in the same ways as those of other IT Sentinel attributes.

The attributes comprising a compound attribute appear as a table of attribute values. The number of rows in this table is specified by an attribute called row. All compound attributes have a row attribute.

Every row of a compound attribute table contains values for the same set of one or more attributes. This set is identified by a numeric index in the first column, then by the name of the compound attribute. The rest of the row contains the values of each attribute in the set. These attributes are identified by a name above each column.

Figure 9-6 Compound Attribute Table Dialog Box

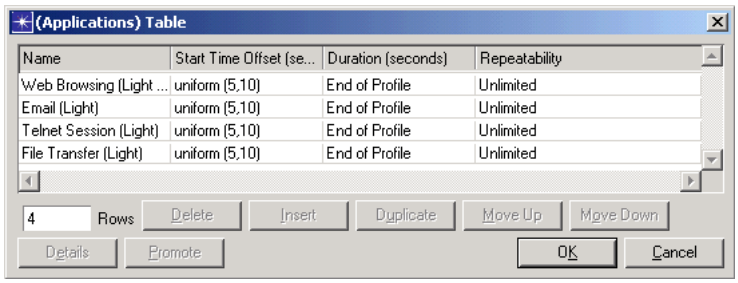

Because a compound attribute incorporates multiple values (for each attribute in the table), its value is represented by a symbol rather than a specific value:

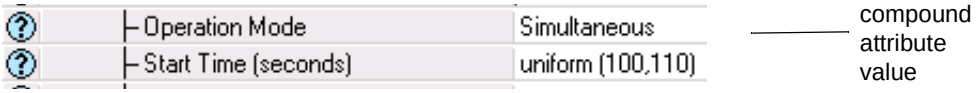

To view the definition of a compound attribute, left-click the “(...)” symbol and choose “Edit” from the pull-down menu.

Defining Compound Attributes

Most of the compound attributes you use will be pre-defined by the model developers. However, unless these compound attributes have been promoted by the developer, you won’t see them. You can create a compound attribute yourself when desired, as an extended attribute.

Procedure 9-14 Defining a Compound Attribute

1. Add the new attribute to the Extended Attributes data table.

2. Select “compound” from the pull-down menu in the Type column of the Extended Attributes data table.

   ➥ The Type column value changes to “compound” and the Default Value column value changes to the compound attribute symbol “(...)”.

Attribute Name	Group	Type	Units	Default Value
operating_system	system	compound		(...)

3 Left-click the Edit Properties button.

➥ The Compound Attributes Properties dialog box appears.

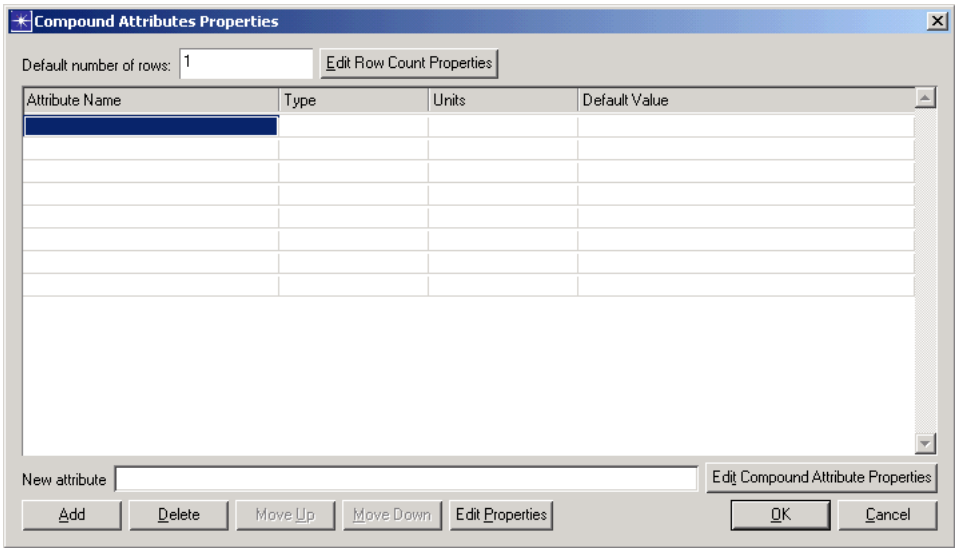

4 Edit the count properties:

4.1 Left-click the Edit Row Count Properties button.

➥ The Attribute Properties dialog box for the count value appears.

4.2 Edit the properties of “count” as necessary, using the procedure described in section Attribute Properties on page ISU-9-18.

4.3 Close the Attribute Properties dialog box when finished.

5 Edit the set of attributes that will be part of each row of the compound attribute table:

5.1 Type the name of an attribute in the New Attribute text entry area, then press <Return>.

5.2 Set the type and default value in the corresponding columns.

5.3 Repeat steps a and b for each additional attribute.

6 If desired, edit other properties of the compound attribute:

6.1 Left-click the Edit Compound Attribute Properties button.

➥ The Attribute Properties dialog box appears. This dialog box allows you to make the attribute properties public, create a symbol map, add comments, and so on. See the procedure for editing attribute properties in section Attribute Properties on page ISU-9-18 for details.

6.2 Close the Attribute Properties dialog box when finished.

7 Close the Compound Attributes Properties and Extended Attributes dialog boxes when finished.

➥ The new compound attribute appears in the Attributes dialog box.

End of Procedure 9-14

Editing Compound Attributes

There are three situations in which you may need to change the values of a compound attribute:

- It belongs to the extended attributes of an object.
- It has been promoted from a lower level of a model.
- It belongs to the parent model of a new derived model.

The procedure for editing a compound attribute is similar in each of these cases.

Procedure 9-15 Editing a Compound Attribute

1 With the compound attribute to be edited visible in the Attributes dialog box, left-click the Extended Attributes button.

➥ The Extended Attributes dialog box appears.

2 Left-click the attribute name, then click in the New Attribute text entry area.

➥ The Compound Attributes Properties dialog box appears.

3 Edit the count, attribute values, and compound attribute properties as necessary.

4 Close all dialog boxes when finished.

End of Procedure 9-15

Customizing Header Information

The Edit Attributes dialog box displays the device type and vendor or product information of the corresponding object, as shown in the following diagram. You can customize this information as follows.

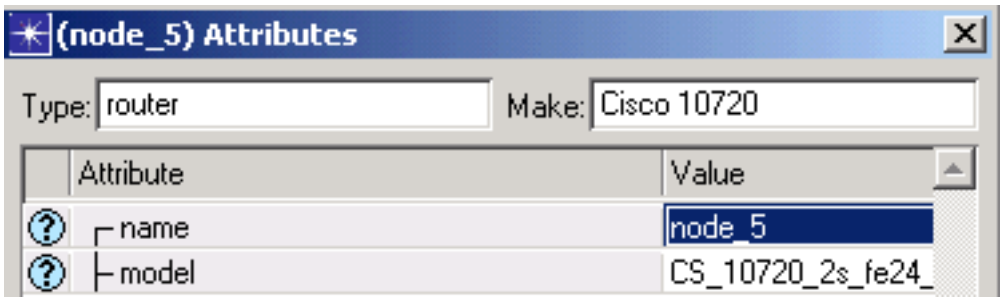

IT Sentinel can read data from the attributes or self-description of an object and display this information in the object's Edit Attributes dialog box. The "obj_dbox_headers" preference specifies the information that appears.

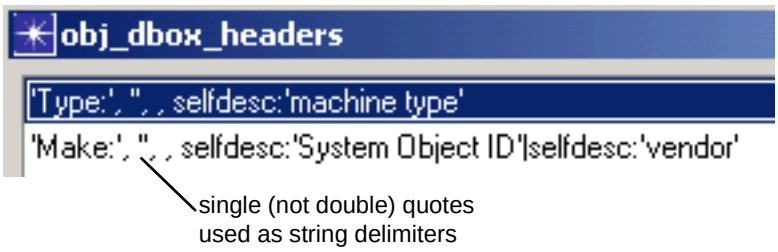

The previous diagram shows the "obj_dbox_headers" table with the default settings. Each line in this table defines a label/value pair: the first line specifies the node type, and the second line specifies the vendor type. (Both these values are specified in the node model's self-description).

Each line in this table specifies four pieces of information, separated by commas:

- <*label_text*>—A label clearly indicating the value displayed, such as `'Type: '`, `'Vendor: '`, or `'Model: '`.
- <*label_tooltip*>—The text that appears in the tooltip when a user rests the cursor over the label.
- <*value_tooltip*>—The text that appears in the tooltip when a user rests the cursor over the label.
- <*value_text*>—The data that appears in the value field. You can specify data from the model's attribute list and self-description for this field.

Note that you can specify multiple display fields; the resulting dialog box displays a maximum of two label/value pairs per line.

The following sections contain notes about how to set these fields.

<label_text> Field

We recommend using the following syntax:

```
<single_quote><label_text><colon><space><single_quote>
```

Example: `'Type:'`

This ensures that a space appears between the label and the field. You must delineate this field using single (not double) quotes.

<label_tooltip> and <value_tooltip> Fields

The following conventions apply for these fields:

- To display the value/label text in a tooltip, specify an empty field—that is, nothing between the delineating commas. The result is that the tooltip displays the corresponding <label_text> or <value_text>.
- To display no tooltip, include two single quotes with nothing between them (`''`) . The result is that no tooltip appears when you rest the cursor over a label or value.

Here's an example:

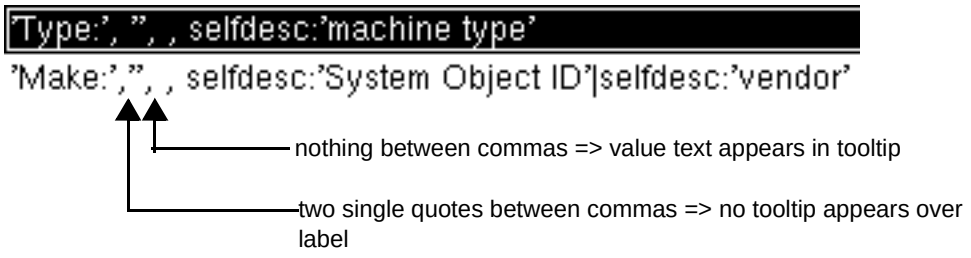

<value_text> Field

This field displays data from a model's attribute settings or its self-description. The syntax is:

```
selfdesc:<self_description_field>

attribute:<attribute_name>
```

You can use the plus sign (+) to show multiple values. For example, to display an object's machine type (router, switch, etc.) and model, separated by a space, you could enter:

```
selfdesc:machine type + ' ' + attribute:model
```

Again, you must use single quotes to delineate characters or strings.

You can use the OR character (|) to create “back-up” values in the event that a particular self-description or attribute value isn’t found. For example, you could enter

```
attribute:<attr1> | attribute:<attr2> | attribute: <attr3>
```

Given this string, the attributes window displays one of the three attributes:

1) IT Sentinel looks for <*attr1*> (attribute 1) in the object’s attribute list. If the list contains <*attr1*>, the attribute window displays this attribute’s value.

2) If the attribute list does not contain <*attr1*>, IT Sentinel looks for <*attr2*> and displays it if possible.

3) If the attribute list does not contain <*attr1*> or <*attr2*>, IT Sentinel looks for <attr3> and displays it if possible.

4) If none of the attributes are present in the list, the header does not appear at all.

End of Procedure 9-15

Attribute Properties Dialog Box

The Attribute Properties dialog box contains a detailed description of an attribute, including its data type, default value, and units. This dialog box is used to view, create, and edit the properties of object attributes. When you access the Attribute Properties dialog box from the Derive New Model dialog box, you can edit attribute properties. When you access it from the Rename/Merge Attributes dialog box, the Attribute Properties dialog box is read-only. See Attribute Properties on page ISU-9-18 for a description of the fields in this dialog box and instructions for editing them.

When you use the Attribute Properties dialog box to edit attribute properties for a derived model, you can control how the attribute properties of the parent model are used in the derived model. Specifically, you can control inheritance of attribute properties and the status of symbols in the symbol map.

Figure 9-7 Attribute Properties Dialog Box

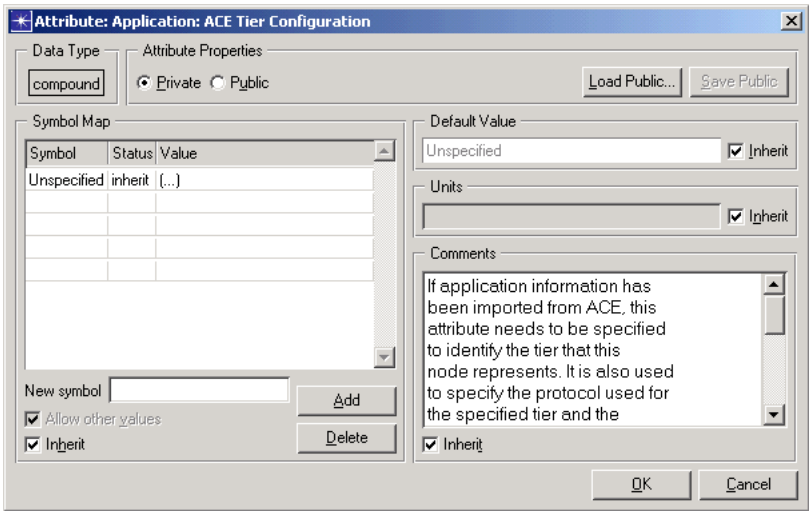

Inheritance

The attributes of derived models automatically inherit the attribute properties of the parent model. IT Sentinel lets you control which properties get inherited via the Inherit check boxes in the Attribute Properties dialog box. These check boxes are only present when editing derived models, not when viewing them. When selected, an Inherit check box indicates that the current model has inherited the properties of its parent model.

An Inherit check box associated with each of the following attribute properties:

- Default value
- Units
- Range
- Comments
- Symbol map

Note—The Inherit check box for a symbol map pertains to the status of the Allow Other Values check box only, not to the symbol map itself.

When you edit a derived model, you can accept the parent model's attribute property by leaving the check box selected (it is selected by default). While it is selected, you cannot edit the attribute property. Attribute properties in child models whose Inherit check boxes are selected are automatically updated when the parent model is modified.

Deselecting the Inherit check box allows you to modify the attribute property and override the parent model's setting. Any child model derived from it will inherit the modified attribute property.

Figure 9-8 Inherit Check Box

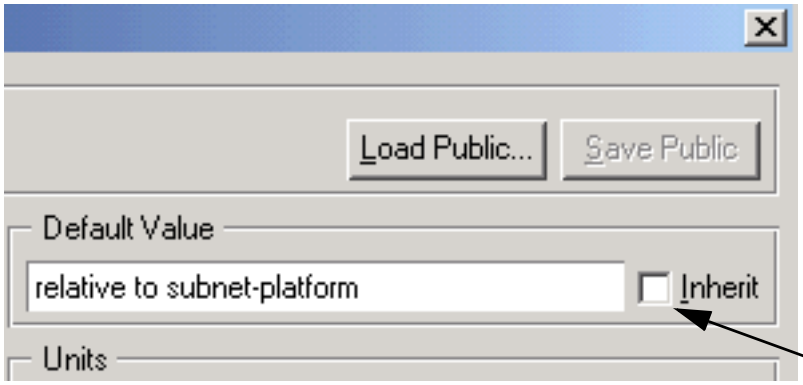

The Inherit check box is deselected, indicating the Default Value attribute property specified here overrides that of the parent model.

Because a derived model is a specialized case of its parent model, you can only override an attribute property with a more restrictive value than that of the parent. For example, to override the "Range" attribute property shown in the above diagram, you would deselect the Inherit check box and then specify a more restrictive range, such as 1 inclusive to 1000 inclusive.

Public attribute properties cannot be inherited from a parent model. Thus, if you make attribute properties public (as described under Sharing Attribute Properties on page ISU-9-20), all Inherit check boxes are deselected automatically. They will remain deselected even if you later change the properties back to private.

Symbol Status

Each symbol in the symbol map of the Attribute Properties dialog box for a derived model has a status that specifies how that symbol will be treated in the derived model. New symbols added while creating or editing derived models always have the status "add". Symbols inherited from a parent model's symbol map can have the status "inherit", "replace", or "suppress". The following table describes each symbol status.

Self-Description Dialog Box

The Self-Description dialog box is used to view or modify a model's self-description. Every link and node model contains a self-description that stores basic information about itself, such as machine type, manufacturer, and interfaces. IT Sentinel uses this information to match models to imported objects during topology imports. (For more information, see the chapter on Importing Network Data.)

By editing the self-descriptions of derived models, you can fine-tune how IT Sentinel selects models during a network import. Link and node self-descriptions have different formats, and the procedures for editing them differ.

A self-description has three basic components: a priority, requirements and assignments:

- Priority—If IT Sentinel finds more than one model for an imported object, it chooses the model with the highest priority number.
- Requirements contain information about a model that must match that of an imported object for the model to be selected.
- Assignments are values assigned to a node model's attributes after it is substituted for an imported object. Assignment values of type "constant" are simply assigned to the node's attribute. Assignment values of type "characteristic" use data from the imported object (such as the number of workstations in a LAN) to set an attribute.

Node Self-Description

A node self-description has two components: a core self-description and one or more port self-descriptions. A core self-description stores basic information about a node model, such as machine type, manufacturer, etc. Port self-descriptions store information about one or more ports supported by the node. Each transmitter-receiver pair must be described, though one link self-description can describe multiple pairs if they have the same description and multiple port numbers.

Procedure 9-16 Editing a Node Self-Description

1 Right-click a node's icon in the Object Palette.

➥ The Model Description dialog box opens.

2 Open the Self-Description Dialog box.

If you are editing an existing model's self-description, press Edit in the Node Model Description dialog box. If you are deriving a new model, click Derive New Model and then click Self-description.

➥ The Node Self-Description dialog box opens.

Figure 9-9 Node Self-Description Dialog Box

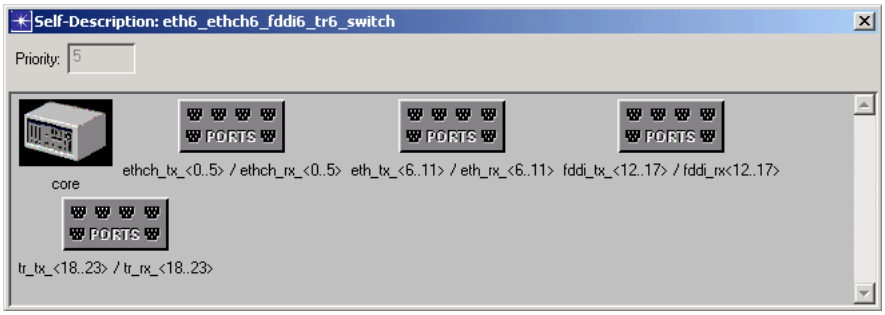

3 If you are deriving a new model, uncheck the Inherit check box.

A derived model inherits its parent's self-description by default, so this box is initially checked and all fields are read-only. You must uncheck this box before you can modify the self-description.

4 Set the priority number in the Priority field.

5 Right-click the core self-description icon.

➥ The Core Self-Description dialog box opens.

Figure 9-10 Core Self-Description Dialog Box

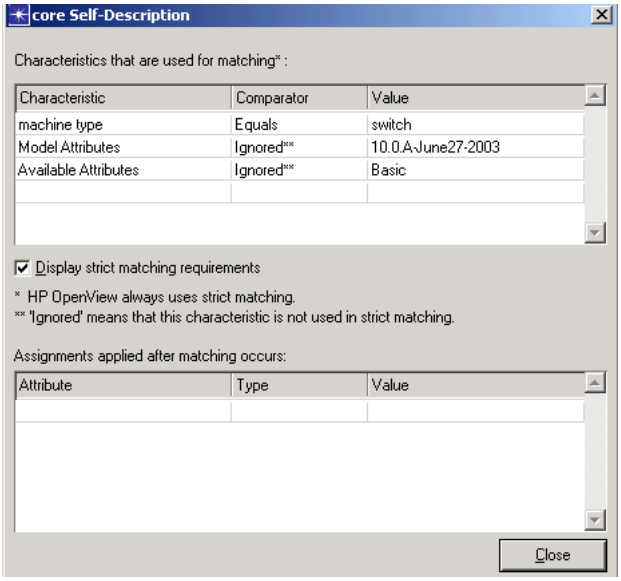

6 Edit the Requirements and Assignments fields as needed and click Close.

7 Right-click on a port self-description icon in the Node Self-Description dialog box.

➥ The Port Self-Description dialog box opens.

Figure 9-11 Port Self-Description

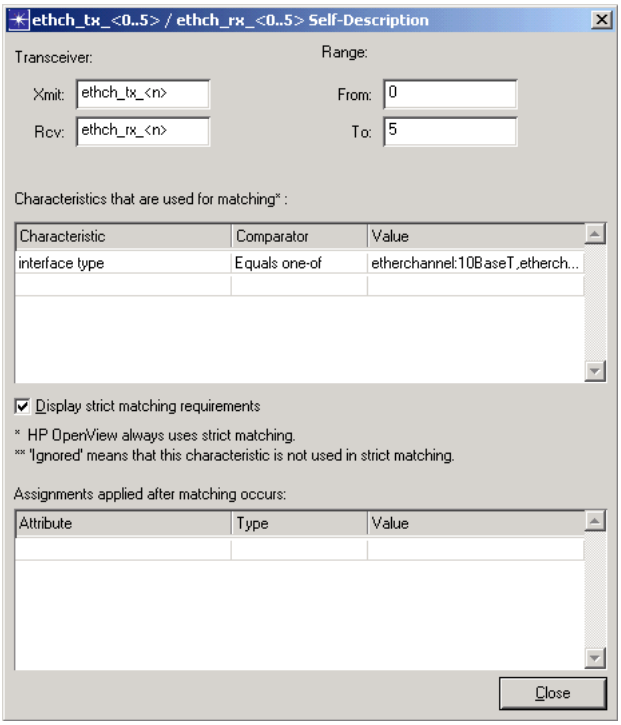

8 Enter the names of a transmitter–receiver pair in the Xmit and Rcv fields. To specify multiple transmitter–receiver pairs with consecutive port numbers, use “<n>” instead of a port number in the name (for example, “hub_tx_<n>_0”) and enter the minimum and maximum port numbers in the Range fields.

9 Specify any requirements and assignments for the node model. Port assignment attribute names can contain the “<n>” wildcard described in the preceding step. Then click Close.

10 Edit, add, or delete any other port self-descriptions, as needed.

11 Click Close in the Node Self-Description dialog box.

End of Procedure 9-16

Link Self-Description Dialog Box

Link self-descriptions have the same elements as node self-descriptions, but are much simpler. One dialog box displays all self–description elements—requirements, assignments, and priority number.

Procedure 9-17 Editing a Link Self-Description

1 Right-click a links’s icon in the Object Palette.

➥ The Model Description dialog box opens.

2 Open the Self-Description dialog box.

If you are editing an existing model's self-description, press Edit. If you are deriving a new model, click Derive New Model and then click Self-description.

➥ The Link Self-Description dialog box opens.

Figure 9-12 Link Self-Description Dialog Box

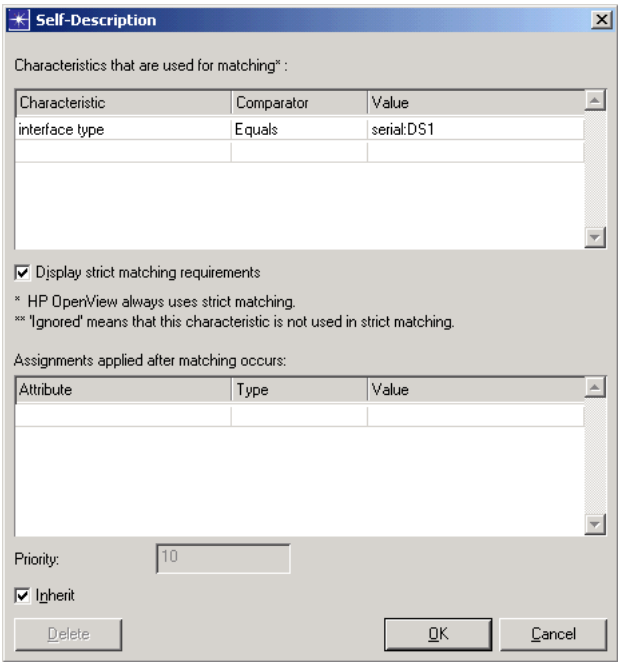

3 If you are deriving a new model, uncheck the Inherit check box.

4 Edit the port self-description as needed, then click Close.

End of Procedure 9-17

10 Importing Network Data

In IT Sentinel, you can import network data using the following methods:

- VNE Server Import on page ISU-10-2
- Device Configuration Imports (DCI) on page ISU-10-15

VNE Server Import

You can create a network model based on information imported from VNE Server. Unlike the other imports described in this chapter, you can import traffic information–not just topology information–when importing VNE Server data. Procedure 10-1 describes the process for importing from VNE Server.

When you choose to import from VNE Server, the following dialog box appears:

Figure 10-1 VNE Server Import Dialog Box–(v11.0 and higher)

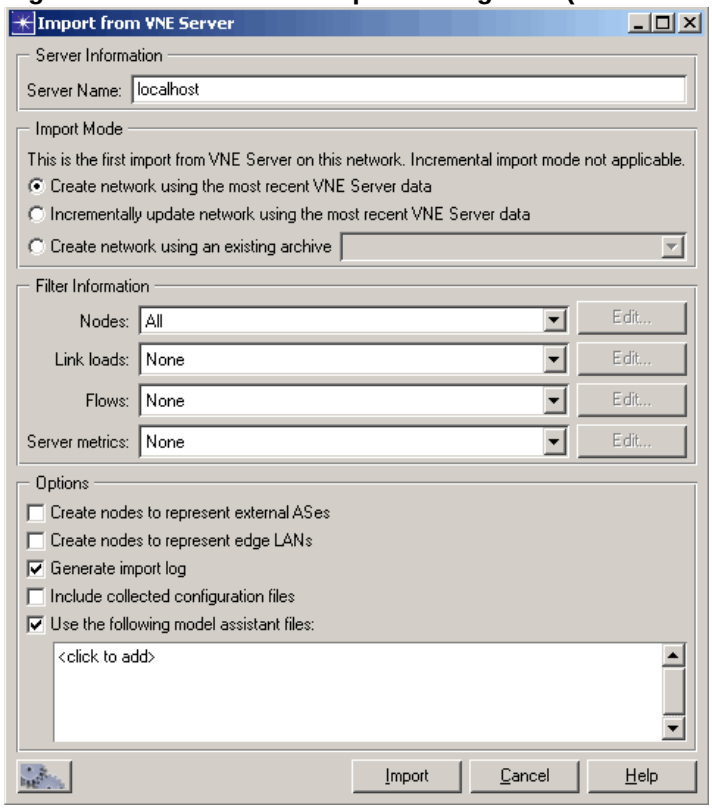

Note—In IT Sentinel versions prior to 11.0, you had to enter the path to the .ior file instead of the server name. If you are using an earlier version of IT Sentinel, please note that you still need to enter the information in the earlier format.

WARNING—If you are using VNE Server versions prior to 3.0, you cannot perform an import by specifying a VNE Server name. You must set the preference vne_import.use_ior_file to TRUE and specify the .ior file to use.

To import the data from VNE Server, you must specify the following information:

- Server Name—Specify the hostname and port for the VNE Server.
  - — If you are running import on the same machine as the VNE Server, use “localhost” as the server name.
  - — If the VNE server is in a different domain, enter the server name as *hostname.domain,* i.e. vneserver.opnet.com.
  - — If you are importing from VNE Server on a non-standard port (i.e. not port 900), enter the server name as *hostname:port* (or *hostname.domain:port*).

Note—You can change the default port used during import by setting the preference vne_import.vne_server_port to the desired port number.

- Import Mode—See Table 10-1 for information on import modes.
- Filter Information—See Table 10-2 for information on filtering data during import.
- Options—Specify optional parameters for import.
  - — Create nodes to represent external ASes—Creates dummy objects to represent missing external BGP (EBGP) neighbors.
  - — Create nodes to represent edge LANs—Creates LAN objects for active router interfaces that are not connected. By connecting these interfaces to LAN objects, no active router interfaces are left unconnected after the import.
  - — Generate import log—Creates a log of actions during the import. You can access the log through the object or workspace pop-up menus after the import.
  - — Include collected configuration files—Gathers the configuration files of the devices to be imported along with all other import data. When this option is used, you can right-click on any imported node for which a configuration file exists and select View Device Configuration Source Data.
  - — Use the following model assistant files:

    Model Assistant Files (optional)—If the “Use the following model assistant files:” check box is selected, you can choose model assistant files that specify certain import-related information (i.e., data rates, connectivity). For more information about Model Assistant Files, refer to Model Assistant on page ITU-6-48.

Specify the import mode you wish to use, based on the following table:

Table 10-1 VNE Server Import Modes

Import Mode	Usage	Description
Create/Replace network using the most recent VNE Server data	Use when starting a new scenario or when you wish to erase all existing information in the current model to replace with VNE Server data.	If checked, the current VNE Server data will be used to create a network model. “Create” will appear if you are in a new scenario; “Replace” will appear if you are in an existing scenario.
Incrementally update network using the most recent VNE Server data	Use when you wish to update the current model with any changes that have occurred in VNE Server since the last import.	If checked, VNE Server will update the current network model with all recorded changes in VNE Server since the original import.
Replace network using an existing archive	Use when you want to replace the current model with information contained in a VNE Server archive file.	If checked, a network model will be created using the archived network specified in the drop-down menu. This menu is populated based on available archives from VNE Server.
End of Table 10-1		

Specify filters to apply to the imported data, as described in Table 10-2.

Table 10-2 VNE Server Import Filters

Filter	Description
Nodes	Select "All" to import all available nodes or choose "Filter" to allow selective import. If "Filter" is selected you can also use the "Edit" button to select which node groupings will be imported. These node groupings are created in VNE Server.
Link Loads	Select "All" to import all available link load data, "None" to filter all link load data, or choose "Filter" to allow selective import. If "Filter" is selected you can also use the "Edit" button to select the time period. Only the data currently available in VNE Server will be selectable.
Flows	Select "All" to import all available link load data, "None" to filter all traffic flow data, or choose "Filter" to allow selective import. If "Filter" is selected you can also use the "Edit" button to select which traffic flow data and for what time window. Only the data currently available in VNE Server will be selectable.
Server Metrics	Select "All" to import all available link load data, "None" to filter all server metric data, or choose "Filter" to allow selective import. If "Filter" is selected you can also use the "Edit" button to specify the data to be filtered. Only the data currently available in VNE Server will be selectable.
End of Table 10-2	

If you filter based on groups and do subsequent incremental imports, the original node grouping that was in VNE Server at the time of the original import is retained. Changes, additions, or deletions from the original node grouping are not considered during incremental imports.

Note—Filters do not apply when performing an incremental import from VNE Server.

Including Device Configuration Files

You can import the device configuration files along with the network model when you perform a new import from VNE Server or when you choose “Replace network using the most recent VNE Server data”. It does not work with import from archive file or with incremental imports. See Figure 10-2 for specific configuration information.

Figure 10-2 Include Collected Configuration Files

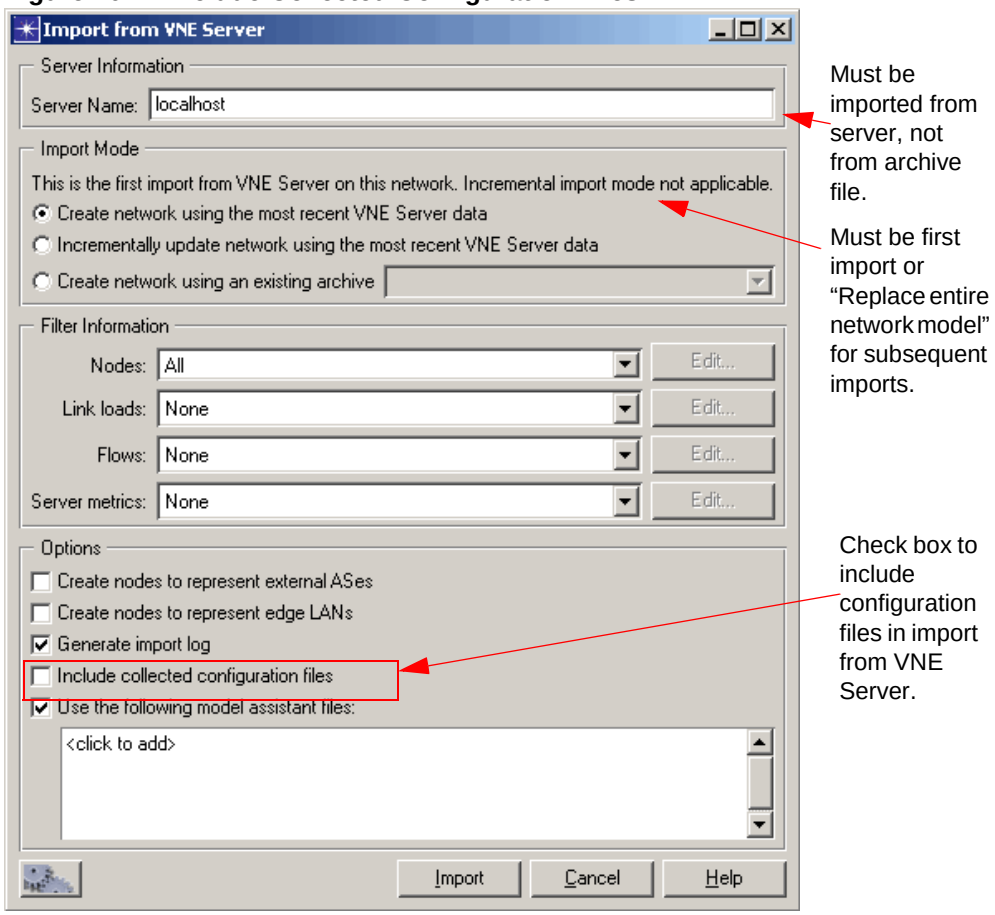

The files that are sent from VNE Server when the “Include collected configuration files” option is selected are config files, “show version”, and “show vlan” data imported into VNE Server using the Device Configuration File Import, Device Version Import, and Device VLAN Database Import adapters, respectively. These files are typically provided by Device Configuration File Collection adapter but can also be copied via the Remote File Collection adapter as is the case with some users.)

Note—With VNE Server version 3.0.1 build 783, Ciscoworks Configuration File Import can also be a source for configuration files, however you MUST have updated model libraries for this feature to work.

Procedure 10-1 Importing a VNE Server Network into IT Sentinel

1. Start VNE Server if it is not already running. Services must be running to import from VNE Server. See the VNE Server *User Guide* for more information.
2. You can import your VNE Server topology in one of two ways:
   - To import into a new scenario: choose Import from VNE Server in the Initial Topology window of the Startup Wizard.
   - To import into an existing scenario: from the Project Editor's Topology menu, choose Import Topology > From VNE Server... (Keep in mind that importing into an existing scenario will overwrite the original contents.)

   ➥ The Import from VNE Server dialog box appears.
3. Specify VNE Server hostname (or the location of the vneserver.ior file is located), and configure import options.

   Note—If there is a firewall between Modeler and the VNE Server, use the .ior file. To do this, set the preference vne_import.use_ior_file to TRUE through the menu path Edit > Preferences or by editing the env_db file with a text editor.
4. Click Import.

   ➥ IT Sentinel imports the topology and traffic (if selected).

End of Procedure 10-1

After the import has completed, an import summary is displayed, as shown in Figure 10-3. The summary contains information about the devices and links imported and created in the model. It also shows the timestamp of the import and in what mode the network was imported.

Figure 10-3 Import Summary Report

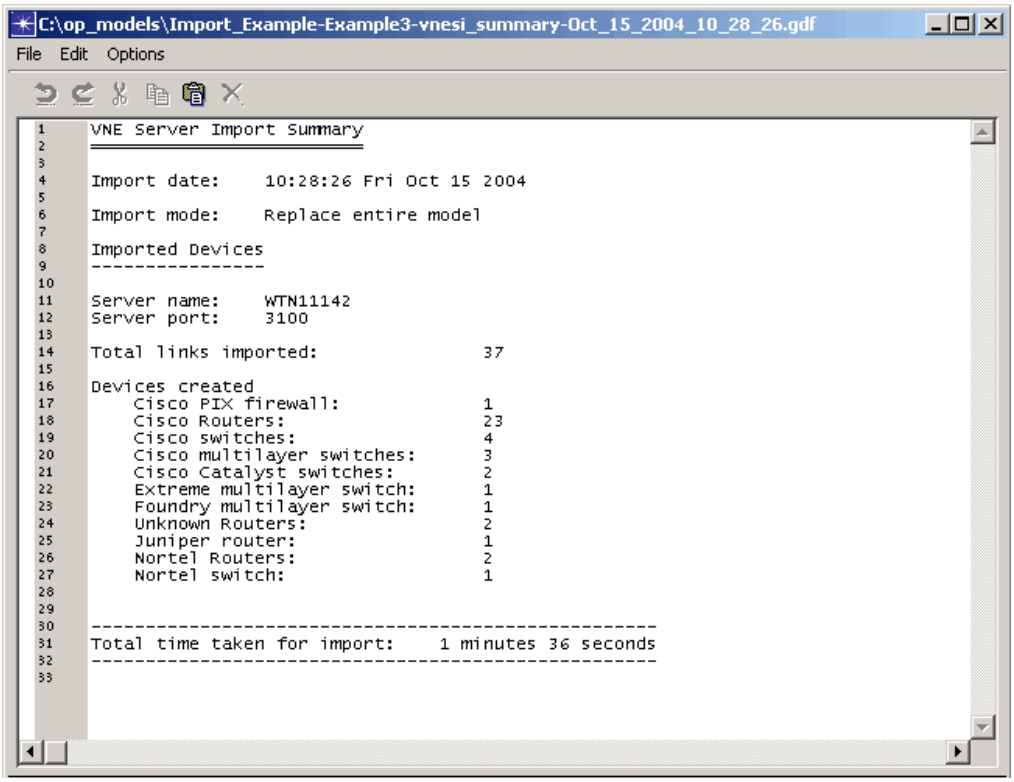

After the Import

One you have imported your network model from VNE Server, you can

- Perform an Incremental Import
- View and Edit Configuration Files
- Use the Virtual Command Line Interface (Virtual CLI)
- Generate a Network Inventory Summary Report
- Troubleshooting the Import

Incremental Import

Once you have imported a network from VNE Server, you can perform incremental imports. With this method, you import only the changes that have occurred since the last import from VNE Server.

WARNING—In order to do an incremental import from VNE Server, you must set persistChanges to TRUE in VNE Server. To do this, open the VNE Server Management Console and navigate to Project Properties > VNES Features. This attribute is set to FALSE by default.

Notice that instead of the warning “This is the first import from VNE Server on this network. Incremental import mode is not applicable”, as shown in Figure 10-1, you are now given the timestamp for the last import from VNE server, shown in Figure 10-4.

Figure 10-4 Incremental Import from VNE Server

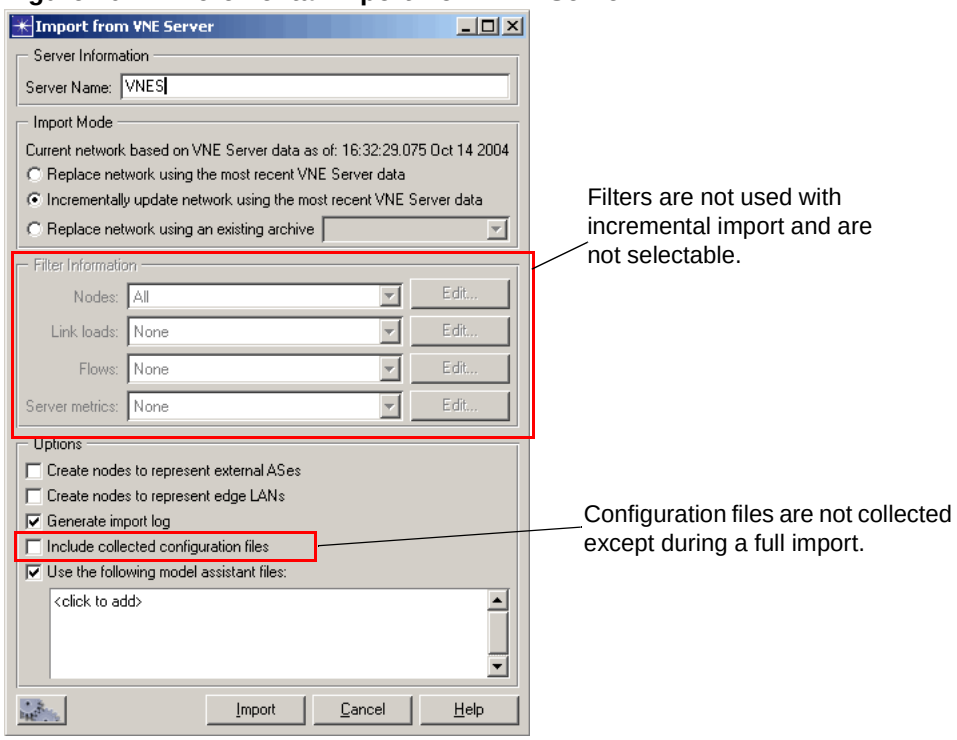

If you select “Incrementally update network model using the most recent VNE Server data”, as shown in Figure 10-4, any changes that occurred in the VNE Server model since the last import will be imported. This is a simpler method of importing only the most recent changes, allowing you to audit the new network model and perform comparison studies.

View and Edit Configuration Files

After an import, you can view the configuration file that import used to create a node in the Project Editor. To do this, select View Device Configuration Source Data from the node object’s pop-up menu.

WARNING—Changing your device configuration source files and then performing an incremental import on the modified network model might give unexpected results. It is important to do incremental imports only to the original network model imported. If you want to make modifications to the model, save the scenario under another name first.

Procedure 10-2 Editing a Configuration File of a Device

1. Right-click on the device in the Project Editor workspace and select View Device Configuration Source Data from the pop-up menu.

   ➥ The configuration file appears in a text edit pad (see Figure 10-24 on page ISU-10-38).

2. Edit the configuration file.

3. Choose File > Save to save your changes.

4. Close the text edit pad.

End of Procedure 10-2

Figure 10-5 Viewing the Device Configuration File After Import

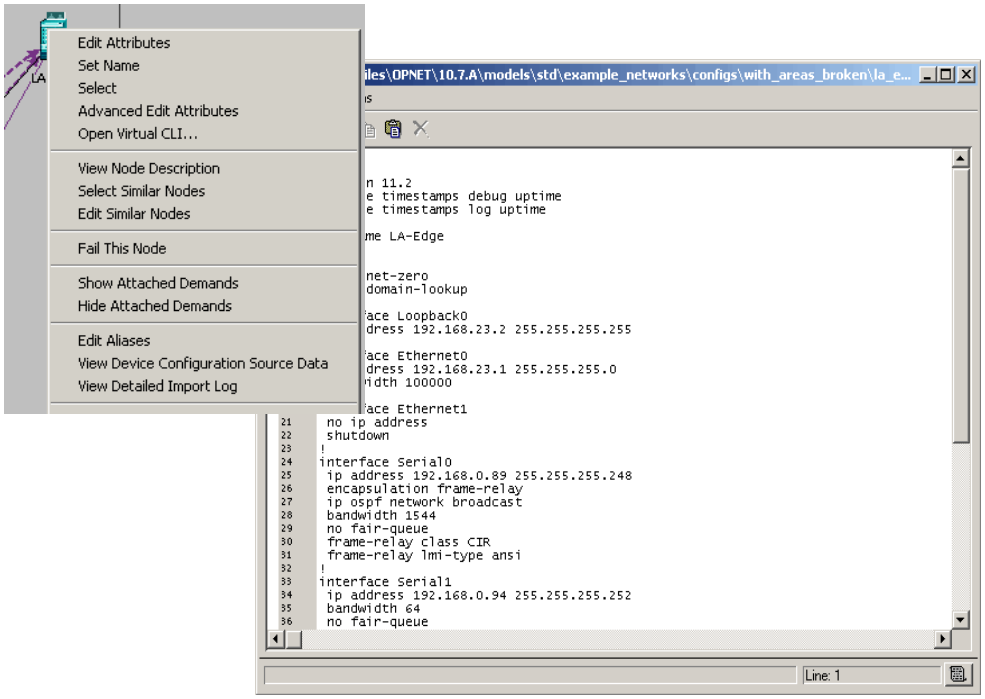

Virtual Command Line Interface

The last section described how to permanently edit your configuration files through a text editor. If you want to edit the devices only for a particular project, you can use Virtual Command Line Interface (Virtual CLI), which allows you to edit Cisco routers and switches within the network model. The changes made with Virtual CLI are specific to the project within which you are working and do not affect the underlying source files.

WARNING—If an incremental import is performed in the project after changes are made through Virtual CLI, those changes are disregarded. It is recommended that you save your original project and scenario and do all modifications in a duplicate scenario, saved under a different name.

The Virtual CLI has the same look and feel as the Cisco CLI, although it supports a sub-set of Cisco commands. Virtual CLI supports familiar Cisco CLI usage including shortcuts (such as typing `ena` instead of `enable`), command help access via "?", and tab completion of a command.

To access the Virtual CLI, right-click on a Cisco router or switch and select Open Virtual CLI... Figure 10-6 shows an example of a Virtual CLI session.

Figure 10-6 Virtual Command Line Interface

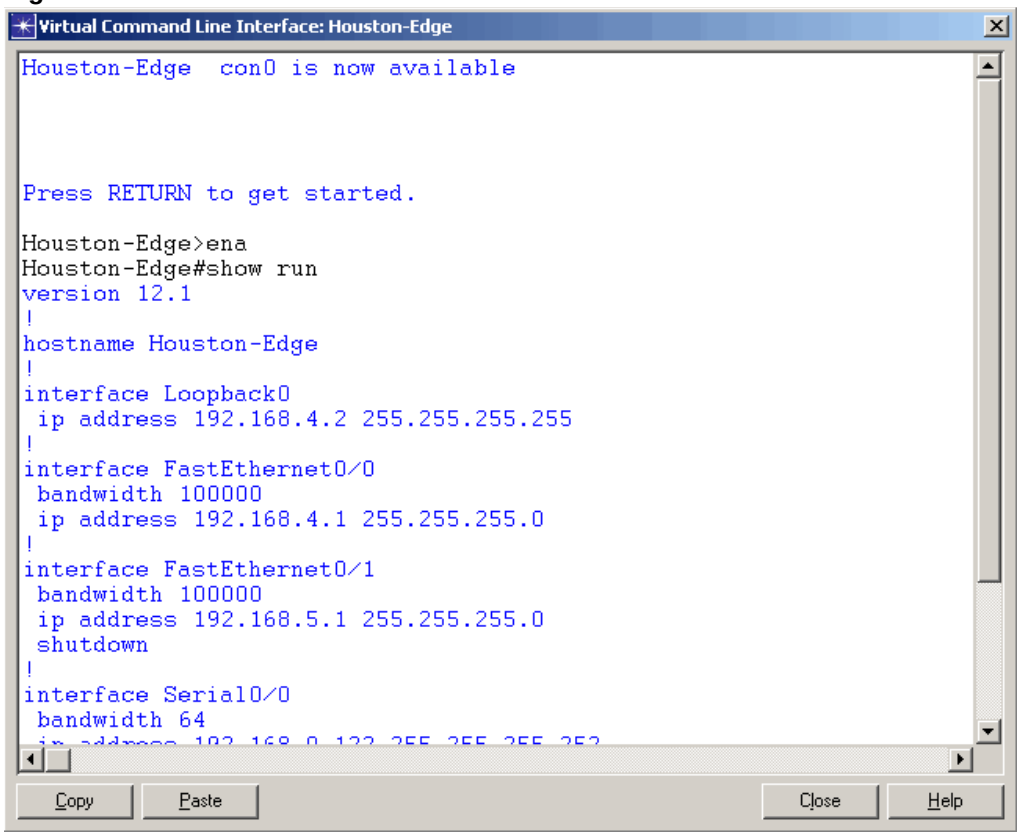

Network Inventory Summary Report

After an import, and on a periodic basis, you may want to see how many devices are in your network and of what type. If you have a rather small imported network model, it is a simple task to visually gather this information. For larger networks, however, it is a time-consuming, error-prone task to take inventory.

The Network Inventory Summary report is available at any time after an import. Running a Network Inventory Summary report after an incremental import allows you to keep your inventory files up-to-date. In Figure 10-7 you can see that there are two ways to generate the report: through the drop-down menus or through the toolbar.

Figure 10-7 Generating a Network Inventory Summary

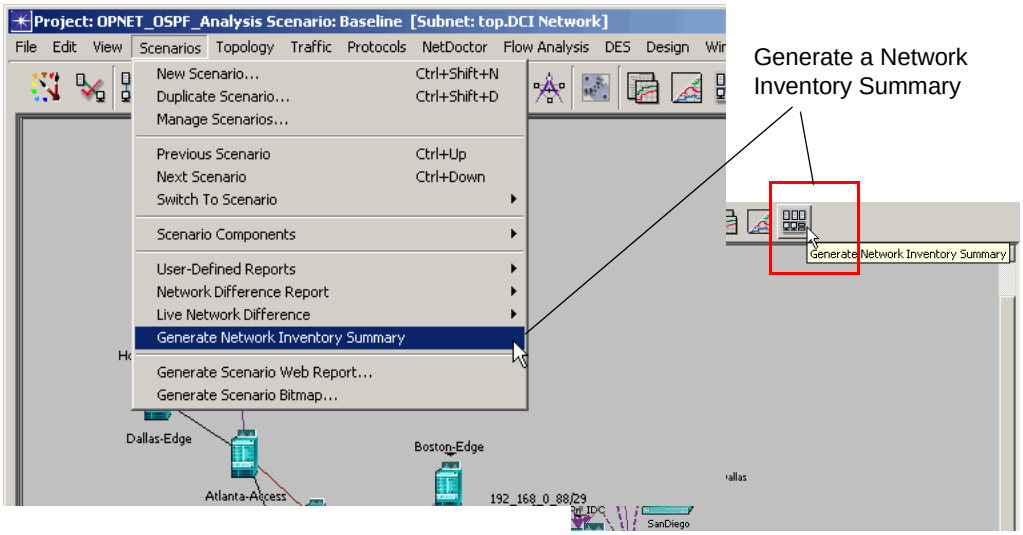

An example of a Network Inventory Summary report is shown in Figure 10-8. The summary includes

- Devices—A total count of devices, and a breakdown by device type
- Vendors—A count by vendor
- Physical Links—A total count of physical links, and a breakdown by type
- Virtual Links—A total count of virtual links, and a breakdown by type
- Other—A total count of other types of devices in the model, including network clouds, application models, configuration utilities, etc.

If a count listing is in blue text, there are further actions you can take. In the example shown in Figure 10-8, certain counts are in blue text. To drill down even further, you can click on any number in blue. In the example below, we can go two levels deep to see the four Cisco WS-C2924M-XL switches listed.

Figure 10-8 Network Inventory Summary

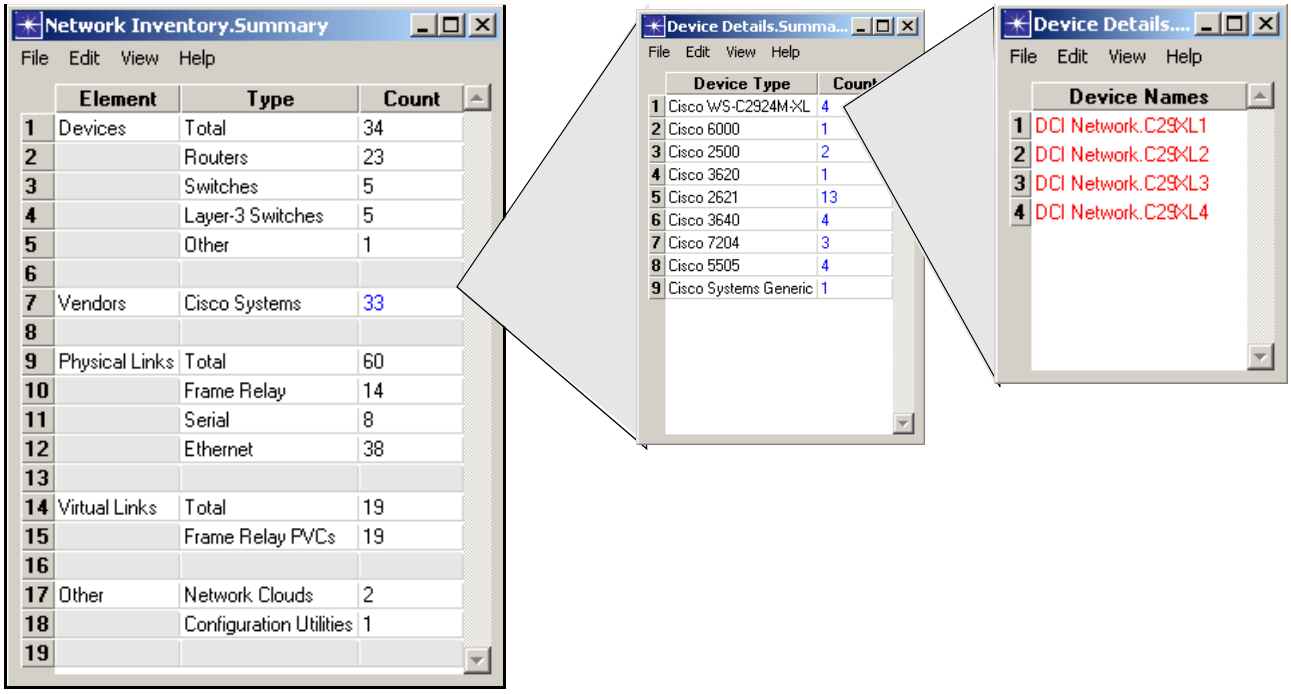

At any point in the summary, you can also export the report to any of the following file types:

- Spreadsheet (for example, MS-Excel)
- Comma-Separated ASCII
- Web Report
- XML

Figure 10-9 Export Network Inventory Summary

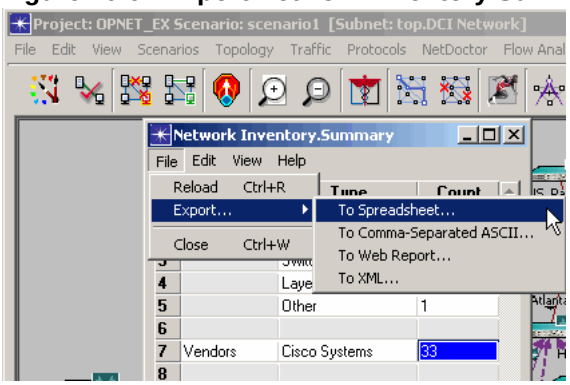

Troubleshooting the Import

Occasionally, you may encounter an unexpected problem during an import or in the results after the import. Some common issues are presented here.

VNE Server Behind Firewall

If there is a situation in which the VNE Server cannot be easily reached (for example, if a firewall separates the VNE Server and Sentinel), you may import from a .ior file. To do this, you must set the preference vne_import.use_ior_file to TRUE. This changes “Server Name:” to “IOR File Location:” in the dialog box.

Links with Mixed Technologies

By default, all simple links with mixed technologies are converted to serial links during import from VNE Server. In clouds with a mixture of at least two of the following technologies—Frame Relay, ATM, and serial—the interfaces are also mapped to serial switches. Interfaces on the connected devices are also changed to serial to prevent a red X from appearing on the link, as though it were unreachable. If you want to disable this default behavior, change the preference vne_import.create_serial_cloud to FALSE.

Failure to Connect to VNE Server During Import

You may encounter a problem connecting to your VNE Server during an import. There are several reasons this can happen.

- VNE Server is not up or its services are not running.
- You did not specify the correct VNE Server hostname or did not specify a fully-qualified hostname and port, if applicable. Refer to the beginning of this section for more information on fully-qualified hostnames.
- If you have the preference set to use the vneserver.ior file instead of a server name, you may have either specified an incorrect .ior file name or the VNE Server services are down.
- You may not have connectivity to the VNE Server. Perform basic network troubleshooting to determine if connectivity exists.
- You may have a firewall between your machine and the VNE Server. See VNE Server Behind Firewall for more information.
- You may have set the preference vne_import.skip_vne_connection to TRUE and have no valid XML file in your <home>/op_admin/tmp directory.

Operational Status “Not Present”

If you perform an import and see a red X on an interface that is operational in the production network, it may be due to a known bug in a Cisco MIB for the interface. When VNE Server performs an SNMP query on the ifOperStatus object or the propPointToPointSerial encapsulation for a High Speed Serial Interface (HSSI) running PPP or propPointToPointSerial encapsulation, the value returned may be “6”, meaning the interface is “not present”. This bug is fixed in Cisco IOS 12.2(14.3), but there are workarounds, as well.

There are two ways you can block MIBs from setting the ifOperStatus. You can do this globally by setting a hierarchical priority scheme that applies to all interfaces or by using import blocker on the adapters for each specific interface.

Device Configuration Imports (DCI)

DCI lets you generate a network model by importing configuration data for

- Juniper routers
- Cisco routers
- Cisco PIX devices
- Cisco Catalyst switches

DCI uses the data in the configuration files to set the IT Sentinel attributes that control routing and switching behavior, including

- Connections, data rates, and interfaces (including PVCs, aggregate interfaces, loopback interfaces, tunnel interfaces, VLANs, and sub-interfaces)
- Protocol information for IP, IP Multicast, BGP, EIGRP, OSPF, RIP (versions 1 and 2), IS-IS, MPLS, IGRP, IPX, Frame Relay, ATM, RSRB, DLSW, NAT, QoS, and VPLS
- ACLs (access control lists), route maps, and Juniper firewall filters
- Static routing, default routing, policy routing, and route redistribution
- VLAN, spanning tree, and bridge configuration information
- Cisco security: AAA, TACACS+, Radius, Kerberos

Supported Commands and Features

During import, OPNET considers and parses device configurations and translates these into model attributes. Currently, not all specific configuration commands and values are supported. For information on the commands and features that are supported, see the OPNET Support Center website.

General Workflow

The general workflow to generate a topology is as follows:

1) Create the device configuration files for the routers, Cisco PIX devices, and switches you want to import.

2) Organize the files into separate directories for each vendor.

3) Import the device configurations into IT Sentinel.

4) Specify missing interface, port, or link information such as data rate and unnumbered interface information. DCI includes an Import Assistant for you to specify this information. To see more about this feature, refer to Using the Import Assistant on page ISU-10-28.

To include supplemental information that is not in the configuration files, you can apply model assistant files during or after the import. You can also do a partial or incremental import to add new devices to the topology or to update devices already in the topology.

After the network is generated and fully configured, you can run simulation studies, using discrete event or flow analysis, or validation studies using NetDoctor.

Preparing for the Import

This section introduces the pre-import workflow, including preferences, caveats, and troubleshooting tips that are helpful for preparation before you import your device files. Paying careful attention when planning for an import can help reduce post-import issues.

There are two basic pre-import steps:

- Create directories that are accessible to IT Sentinel for the configuration files.
- Obtain configuration files for the devices you want to import.

Configuration File Directories

All the configuration files for a vendor must be in the same directory. Before performing an import, copy the configuration files into separate vendor directories: one directory for all Cisco files and another for Juniper files.

Creating the Configuration Files

Before importing configuration files into IT Sentinel, you must first generate the files from the configuration information on each device. The configuration files are created by running a few commands, which are listed in Table 10-3, on the router or switch and copying the output to a text file, as shown in Figure 10-10.

Table 10-3 Show Commands Used by DCI

Command	Importance	Usefulness
show running-config	Required for Cisco routers and PIX devices	The output from this command carries the basic information about the Cisco router, PIX firewall, or IOS device, interfaces, and configuration.
show configuration	Required for Juniper routers	The output from this command carries the basic information about the Juniper device, interfaces, and configuration.
show config all	Required for Cisco Catalyst Switches	The output from this command carries the basic information about the Cisco Catalyst switch configuration.
show version	Optional	DCI uses the output from this command to determine the type of device from which it was collected.
show cdp neighbors detail	Required for Cisco switches; Optional for Cisco routers	DCI uses the output from this command to determine connectivity between neighboring devices.
show vlan	Required for Cisco switches running IOS.	DCI uses the output from this command to determine the VLANs configured on the switch.
show frame-relay map	Required to replicate static Frame Relay PVCs in model	DCI uses the output from this command to create Frame Relay PVCs in the network model.
show atm map	Required to replicate static ATM PVCs in model	DCI uses the output from this command to create ATM pvcs in the network model.
End of Table 10-3		

Figure 10-10 Creating Device Configuration Files

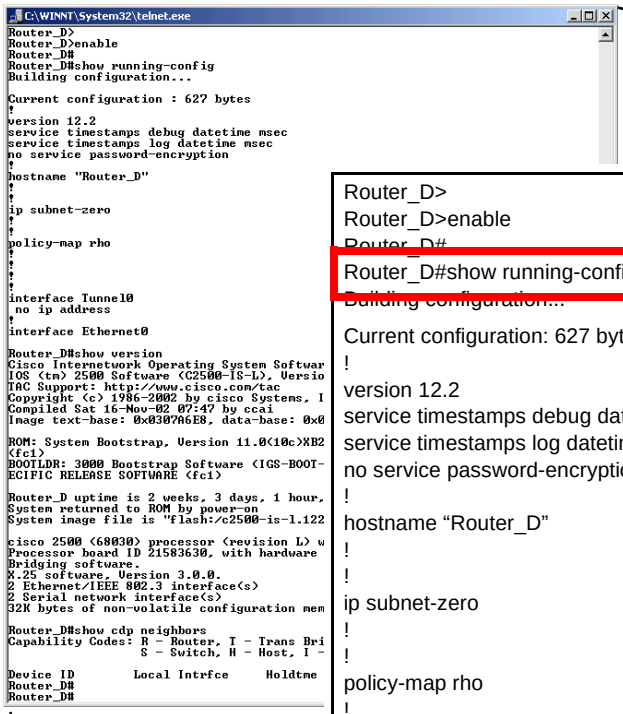

Creating Device Models (Optional)

During the import, DCI searches the model directories for the device model that best matches each configuration file. Typically, this model has more interfaces or ports than the actual device. You can create a set of device models with the same number and type of interfaces as the devices you are importing. Doing this ensures that the models chosen during import match the actual routers or switches exactly.

You can use the Device Creator to create these models, as described in . This step is optional, but it gives you the greatest control over the devices chosen by DCI to represent the imported devices.

Importing Multi-Layer Switch Configurations

In preparation for any device configuration import, you must be sure to gather all the device configuration files necessary to build an accurate topology. It is default DCI behavior to skip a device for which a partial configuration is found. For import of a Layer-3 switch running hybrid IOS (CatOS and IOS), two device configuration files are required. For the import of a Layer-3 switch running Native IOS, only a single configuration file is required.

Figure 10-11 Configuration Files Required for Multi-Layer Switch

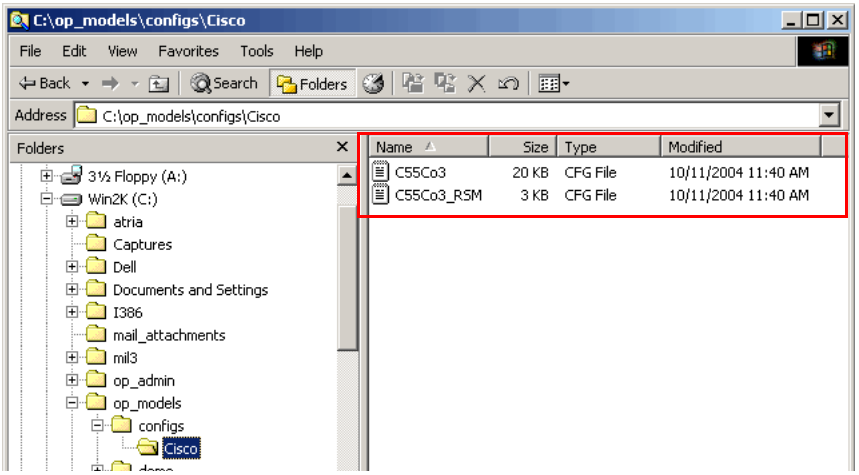

After import, the files create one device in the network model, as shown in Figure 10-12.

Figure 10-12 Import Summary - Multi-Layer Switch

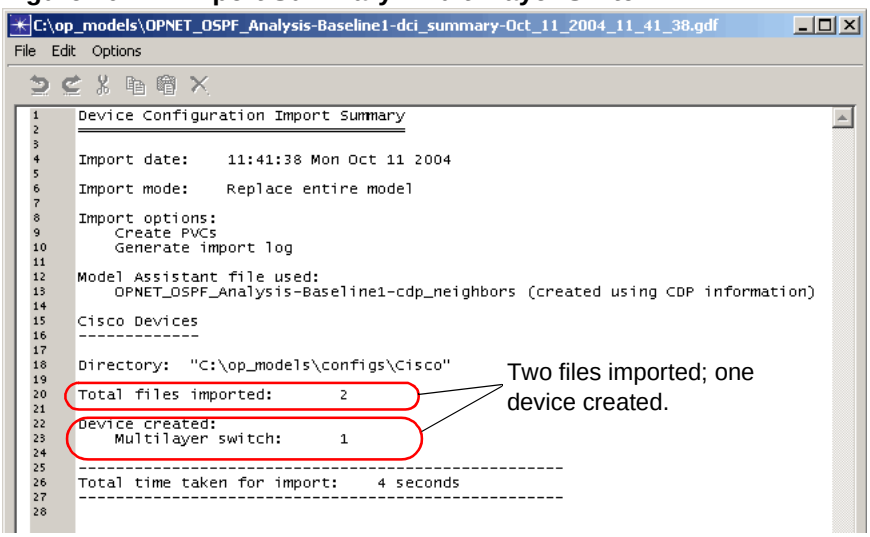

When you want to view the source files for a multi-layer switch running hybrid IOS, the menu choices include both files. You can view source files for both the Layer-2 portion of the device and the Layer-3 portion. Similarly, you can view the import log generated by the import of each configuration file, as shown in Figure 10-13, yet the two files make one node in the model:

Figure 10-13 Source Files and Import Log - Multi-Layer Switch

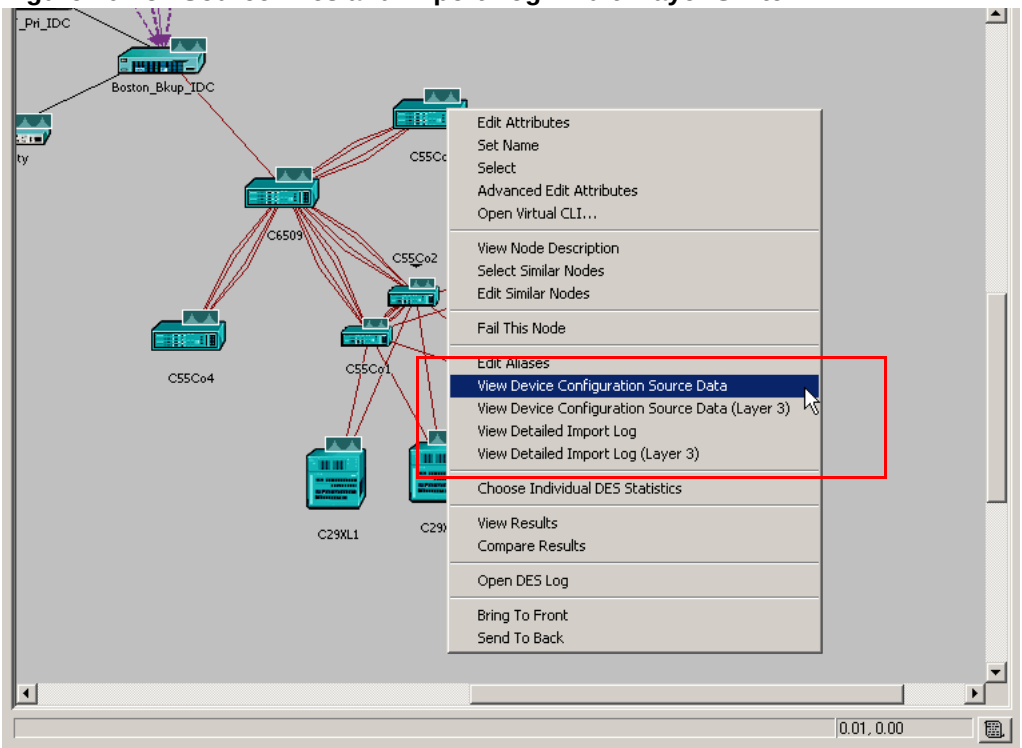

If a multi-layer switch configuration file is examined during import and found to have an active (not shutdown) VLAN interface for which no corresponding switch configuration file is located, however, DCI will skip the import of the entire device with an error message: “File name: *my-file-name.txt* File skipped because Layer 2 configuration file is not available for the multi-layer switch ‘*my-host-name*’.”

Note—Device import preferences are not visible until you have performed at least one device configuration import. If you have not done an import, you can begin an import and cancel out. The preferences then become available.

Restricting Import to Layer-3 Devices

As of release 11.0, you can set a preference to import Layer-3 devices only. When this preference is enabled, DCI will include these devices types in the imported topology:

- Routers/Layer-3 modules (IOS/JUNOS)
- Router-Switches (Hybrid OS/Native IOS)
- PIX Firewalls

It will not include

- Catalyst Switches (CatOS)
- IOS Switches

Note—Router-Switches will only have the Layer-3 components imported when this preference is enabled.

Troubleshooting the Import

The initial network model is built from the imported device files. If adequate planning is not done before the import, the network model might not appear as expected. It is important to understand these issues prior to performing an import, so you may avoid surprises in the resulting topology. After the import, you may notice some of the following issues:

- Missing Devices—Expected devices not showing up in model
- Phantom Devices—Unexpected devices appear in model
- Differences in Connectivity—Physical connectivity is not as expected
- Differences in Devices—Devices are not as expected
- Post-Import Caveats

Missing Devices

You may notice that some expected devices are missing from the network model after an import. Missing devices can be caused by the following expected behavior:

- If a configuration file is not provided for a device, the device is not imported into the network model. Be sure to include all device configurations in the appropriate directory.
- If the `show version` command is not executed for a device, it may not appear as the correct device type in the network model. A device for which there is no `show version` output may appear as a vendor generic device.
- If a VLAN interface is encountered during import, DCI will expect to have a corresponding switch configuration file. If none is found, DCI will skip the device with that interface with an error message. For more information on changing this behavior, see Importing Multi-Layer Switch Configurations on page ISU-10-19.
- If you have a multi-layer switch but did not include the output of `show cdp neighbors detail`, you may have two separate devices appearing in the network model (or neither, as mentioned in the previous bullet). If you choose not to re-import with the cdp information appended, you can create a model assistant file to couple the two devices, using Procedure 10-3. When you perform your next import, be sure to specify the model assistant file you created. The multi-layer switch will appear as one device.

Procedure 10-3 Create a Model Assistant File to Join Multi-Layer Switch

1. Select Topology > Model Assistant > Edit File...

   ➥ The Edit Model Assistant dialog box opens.

2. Click on the Router/Switch Coupling tab.

   Note—You may need to scroll to the right to see this tab. Use the directional arrow in the upper right corner.

3. Click in the Rows box, in the lower left corner, and increment the number by one to add a row.

   Click on Specify... under Router Name, and enter the name of the Layer-3 device.

   Note—This is the name of the RSM/MSFC/RSFC.

4. Click on Specify... under Switch Name, and enter the name of the Layer-2 device.

5. Save the Model Assistant file for use with future imports.

End of Procedure 10-3

Phantom Devices

After an import, the network model may contain devices that you did not include in the configuration data. These inferred devices are created when two or more routers have interfaces to the same IP subnet. In this case, DCI uses a generic cloud model to connect the interfaces. The cloud model used to represent the IP subnet corresponds to the type of interfaces being connected and each cloud node is given the name of the IP subnet it represents.

In the following example network, DCI inferred and created an Ethernet switch between three router interfaces. This occurred because three configuration files specified interfaces that connect to the same IP subnet. All other interface data specified one-to-one connections, so DCI did not create additional switching devices.

Figure 10-14 IP Interface Connections in a Generated Topology

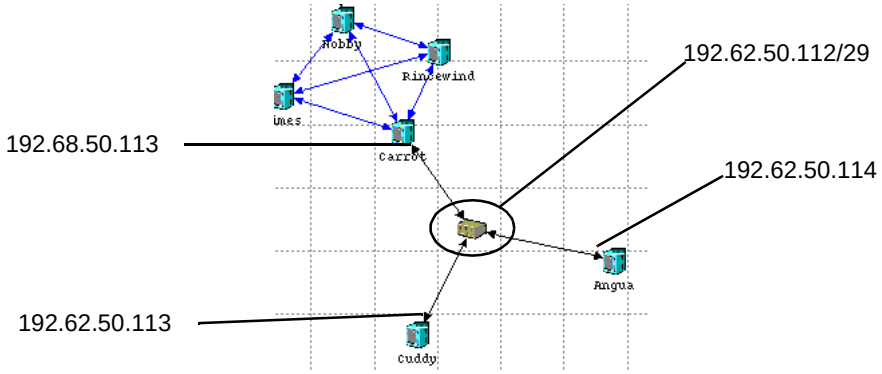

Differences in Connectivity

DCI depends on certain information in the configuration files to make decisions about link names, connections, link types, and link speeds between end-points. The information in this section describes how the connectivity decisions are made.

- When importing router configuration data, DCI relies on the IP, Frame Relay or ATM information in your configuration files to configure router connections.
- By default, DCI names all interfaces based on the names found in the configuration files. If an inferred router model is used to represent the imported device and has extra interfaces, the interfaces are shutdown by default and have generic names in the format “IF*n*” (IF0, IF1, and so on).
- Interfaces that are part of an aggregate interface, such as EtherChannel, are connected based on the configuration of the aggregate interface.
- DCI colors each link according to its underlying technology (Ethernet, Frame Relay, ATM, etc.) and auto-assigns a name that indicates both interfaces to which it is connected.
- DCI treats all DPT/SRP interfaces as serial (PPP) interfaces.

- If the import process cannot determine an interface data rate from a configuration file, and you do not specify a data rate in the Import Assistant, DCI assigns default rates of T1 for serial interfaces and OC3 for ATM interfaces. These default rates can be changed
- Frame Relay and ATM interfaces that are connected through a homogeneous Layer-2 network are imported as Frame Relay and ATM interfaces. However, if these interfaces are connected through a heterogeneous Layer-2 network (such as a mixture of technologies like ATM, Ethernet, Frame Relay within the same IP subnetwork), the interfaces are imported as serial interfaces to allow for connectivity.
- For Frame Relay PVCs, the output from the **show frame-relay map** command will override PVCs configured in interface mode using the command **frame-relay map ip**.
- For ATM PVCs, the output from the **show atm map** command will override PVCs configured in interface mode using the command **pvc** along with the **protocol** command.

Differences in Devices

DCI searches the model directories during the import for the device model that best matches each configuration file. If a matching device is not found, you may discover that the representative device does not suit your needs.

For example, the default device for a Cisco 12012 has 10 ATM ports, 8 Fast Ethernet ports, and so on. Suppose your specific router has no ATM ports but is loaded with Gigabit Ethernet and Fast Ethernet ports. Rather than showing too many ATM ports (in a shutdown status), it is better to define your unique devices in advance of the import.

Adding devices is described in Creating Device Models (Optional) on page ISU-10-19.

Post-Import Caveats

If you plan to import traffic into the network model, there are some things to bear in mind with regard to expected DCI behaviors:

- DCI creates a list of aliases for each router that includes all IP addresses for which that router is responsible. With the exception of switches that have a Layer-3 card, no aliases are created for Layer-2 objects such as switches and clouds.
- By default, the import process does not create any traffic source or destination nodes (such as workstations or LANs). You can create LAN nodes at the edge of the network model by enabling the Create Edge LANs option in the Import Device Configurations dialog box. See Performing an Import with DCI on page ISU-10-25 for details. This option automatically creates Ethernet, Frame Relay, serial, ATM LANs, FDDI, or Token Ring.

- Certain technologies and interface types are not currently supported in DCI:
  - — Dual MSFCs—Only the active router card is imported. If both router cards are active, DCI will import the first one parsed.
  - — Layer-2 Aggregate Interfaces—Layer-2 aggregate interfaces are only supported for Ethernet and Multilink PPP. For other technologies, such as ATM, that have aggregated interfaces, only the physical interfaces are imported. This results in devices connected at the physical interface with no logical aggregation.
  - — Peer interfaces or ports with mismatching technologies may not be connected.
  - — PIX devices may not have the correct chassis type, as the `show version` command is not supported for PIX at this time.
  - — Serial, Frame Relay, and ATM interfaces are automatically converted to Serial and are connected, if they coexist in the same subnet.

Performing an Import with DCI

When you have completed your preparations for an import, you can continue with an import, as described in Procedure 10-4.

Note—Be sure to check the box to “Use Startup Wizard when creating new scenarios” when you create a new project or scenario, as shown in Figure 10-15.

Figure 10-15 Enable Startup Wizard

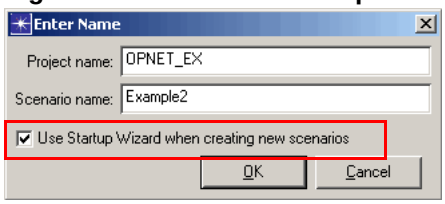

Procedure 10-4 Importing a Topology from Device Configuration Files

1. Open the Import Device Configurations dialog box. You can open this dialog box in one of two ways:
   - To import into a new scenario, choose Import From Device Configurations in the Initial Topology window of the Startup Wizard.

   Note—Be sure to check **Use Startup Wizard when creating new scenarios**, or the Startup Wizard will not appear.

Figure 10-16 Enable Startup Wizard

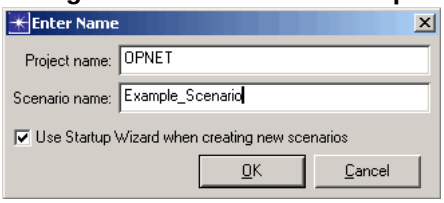

- To import into an existing scenario, choose Topology > Import Topology > From Device Configurations...

Either action opens the Import Device Configurations dialog box, shown in Figure 10-17.

Figure 10-17 Import Device Configurations Dialog Box

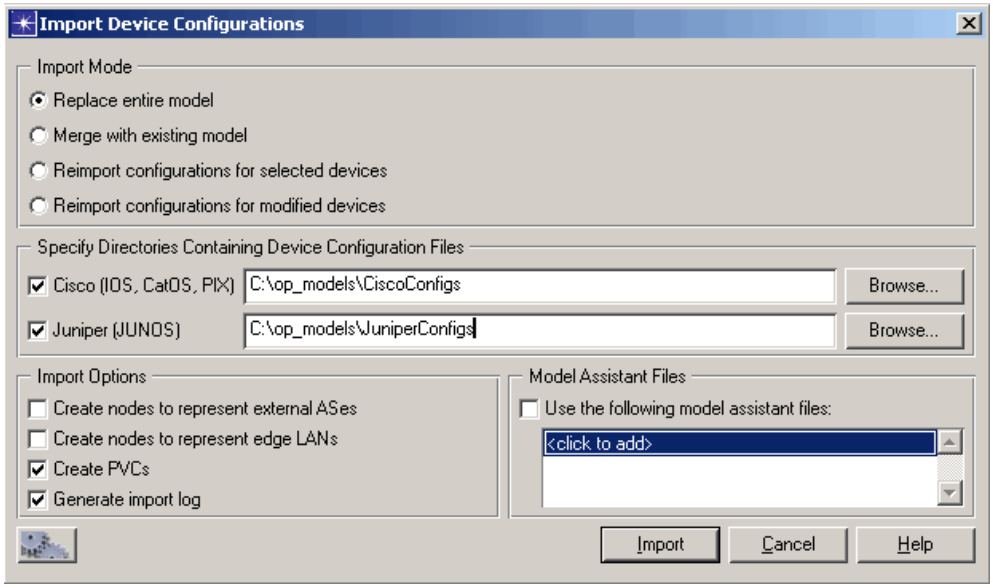

2 Select one of the following import modes.

- Replace Entire Model—Use this option when creating a new project or scenario, or when you want to overwrite the existing topology.
- Merge with Existing Model—Use this option to add devices to the existing topology. If some of the devices are already in the topology, these devices are replaced (that is, reimported) using the configuration files of this import. DCI matches the devices based on the hostname of the node.
- Reimport Configurations for Selected Devices—Use this option to update the configuration files of a few devices in the existing topology. The devices must already be selected in the workspace. If the devices you want to update are not currently selected, cancel this operation, select the devices, and start over.

  In this mode, you do not need to specify the location of the configuration files in step 3, DCI uses the file specified in the Import Information attribute on the device. By default, this is the location specified in the original import.

- Reimport Configuration for Modified Devices—When you select this option, DCI will reimport configuration files for any devices that have been modified since the last import, based on comparison of the timestamp of the last import and the current time.

3 For each type of configuration file from which you want to import, use the following steps to specify where the configuration files are located.

3.1 Select the check box next to the vendor configuration file type.

3.2 Click Browse to locate the directory that contains the configuration files.

3.3 In the Browse for Folder dialog box, select the folder that has the configuration files for this vendor and click OK.

Make sure the other vendors (that is, the ones you do not want to import from) are not selected.

4 Select import options. The options are described in Table 10-4:

Table 10-4 Device Configuration Import Options

Import Option	Description
Create nodes to represent external ASes.	Reads external BGP peering, and if there is no external partner, IT Sentinel creates another AS to connect to. When this option is checked, DCI imports the BGP-RIB information for EBGP connections. The RIB-in data of the border router of the local AS is used as the RIB-out on the external AS node. The RIB-out files are automatically generated and specified on the AS nodes.
Create nodes to represent edge LANs.	Creates LAN nodes at the edge of the network topology. LANs are created on device edges if an IP interface is active but not connected to another device in the import.
Create PVCs	Reads static PVC mapping from configuration and builds the mesh accordingly. If there are no static maps, DCI will create a full mesh of PVCs across all interfaces in the same IP subnet. When this option is not selected, DCI does not create PVCs during the import—even when static mapping is in the imported device configuration. NOTE: This option is new for release 11.0. Previous releases had option “Create Full Mesh PVCs,” which did not read static mapping for PVCs.
Generate Import Log	Creates a log that lists the commands that were skipped (not considered) during the import. The log also contains diagnostic messages generated by the import.
End of Table 10-4	

5 Select or deselect the “Use the following model assistant files” checkbox. If this box is deselected, DCI will not consider any model assistant files specified in step 6. DCI will still use CDP information, however, to infer Layer-2 connectivity and router-switch coupling. If this box is selected (default), DCI will consider the model assistant files specified by the user in step 6.

6 Specify one or more model assistant files. These files can be model assistant files, which have a .ma extension, or import assistance files from releases prior to 10.0, which have a .ia extension. See information about model assistant files.

Note—As of release 11.0, you may select or deselect a checkbox that tells DCI whether or not to consider model assistant files when performing an import. By default the box is selected.

7 If you have a Sentinel or Automation license, and want to save your import settings to an automation file, click the Save Settings for Automation button and specify a name for the automation file.

8 Click Import.

➥ IT Sentinel generates a network topology and configures routing and switching attributes on the nodes according to the information in the configuration files. If additional information is needed to complete the import, a dialog box appears prompting you to open the Import Assistant. See Using the Import Assistant, for more information.

End of Procedure 10-4

The configuration files do not always contain all of the information that can be included in a network model. Although you can manually specify additional information after the import, IT Sentinel has two assistants that can automate the process.

- Import Assistant—When the configuration files do not have some of the information that DCI needs to complete the topology, the Import Assistant appears after the import. See Using the Import Assistant.
- Model Assistant—Unlike the Import Assistant, which prompts you for required information, Model Assistant files let you specify additional information about one, several, or all of the objects in a network. See Using Model Assistant Files with Imported Topologies on page ISU-10-34.

Using the Import Assistant

Sometimes the configuration files do not have all the information that DCI needs to fully configure the device models. When this happens, you need to supply the missing information after the import. The missing information that can be specified in the Import Assistant falls into the following categories:

- IP-unnumbered interfaces—DCI knows that the interface or port is connected to another but is unable to determine the identity of the other interface or port.
- Unspecified data rates—DCI connects two interfaces or ports but cannot determine the data rate of the connection.

When additional import information is needed, the following dialog box appears:

Figure 10-18 Device Configuration Import: Further Information Required

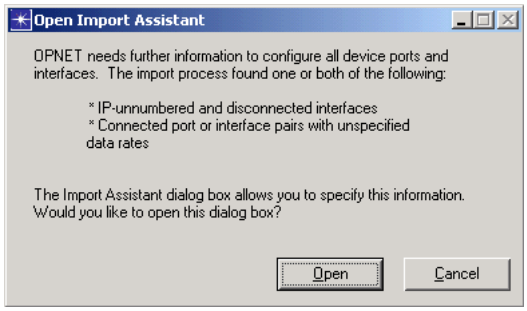

At this point, you can specify information for IP unnumbered interfaces and interfaces with unspecified data rates. Using the Show pull-down menu (top left), you can display either category of interfaces with missing information or all.

Figure 10-19 The Import Assistant

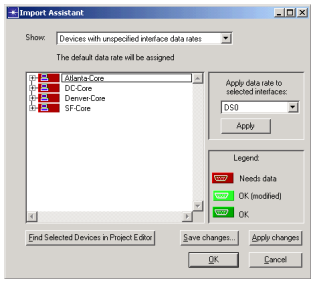

Show routers with unspecified interface data rates

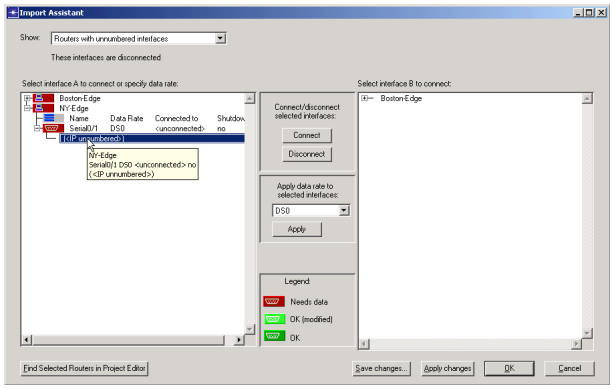

Show routers with unnumbered interfaces

Connecting Unnumbered Interfaces If there are any IP-unnumbered interfaces in the network model, the Import Assistant opens with the Show menu set to Routers with unnumbered interfaces. All such routers (and their unnumbered interfaces) appear with red icons in the treeview. Follow Procedure 10-5 to connect the unnumbered interfaces.

WARNING—This procedure applies only to interface connections that were missing in the configuration data; you cannot use this method to change connections that were set during the import (which appear with dark green icons in the Import Assistant).

When you have connected all unnumbered interfaces, all of the formerly red icons will be light green, indicating that the interface has been modified and is now properly connected.

Note—You can avoid the occurrence of IP-unnumbered interfaces by ensuring that you have the output of `show cdp neighbors detail` command in your configuration files.

Procedure 10-5 Connecting Unnumbered Interfaces

1 Make sure that the Show pull-down menu is set to “Routers with unnumbered interfaces”.

➥ All routers with IP-unnumbered and disconnected interfaces appear in the Interface A pane.

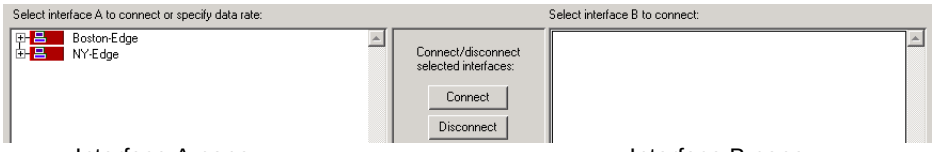

Interface A pane Interface B pane

2 In the Interface A pane, select an unnumbered interface. Note that clicking on the plus sign (Windows) or arrow (Solaris) next to the router icon will expand the view of the router to include its name and interfaces.

➥ A list of routers with eligible interfaces appears in the Interface B pane.

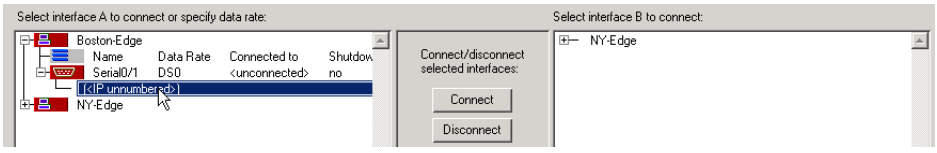

3 In the Interface B pane, select the interface you want to connect to the interface selected in the Interface A pane.

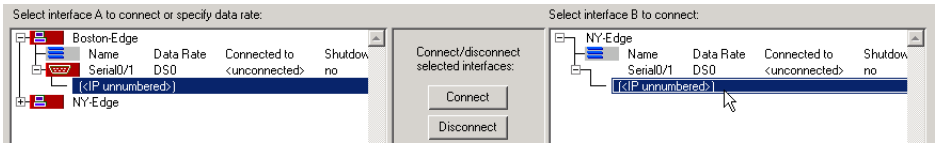

4 Click Connect.

➥ The two interfaces are connected. Interface B disappears from the right pane; the icons for interface A change from red (Needs data) to light green (OK, modified).

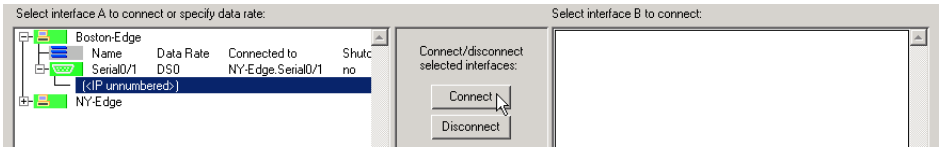

End of Procedure 10-5

Specifying Data Rates You can set or change the data rate for any connected router interface in your network. When you set a data rate, the change affects both connected interfaces.

Procedure 10-6 Specifying the Data Rate of an Interface

1 Set the Show pull-down menu to “Routers with \

2 Data Rates” or “All routers”.

➥ A list of all routers in the network appears. The routers that have interfaces with unspecified data rates appear with red icons.

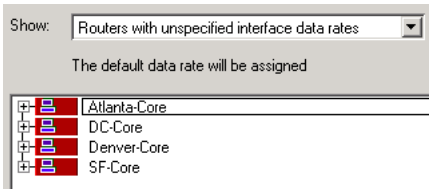

3 Select the interface whose data rate you want to change. You can select multiple interfaces by holding down the Control key and clicking on multiple interfaces. The data rate indicated for an interface is the default data rate that will be assigned if you do not specify the actual data rate.

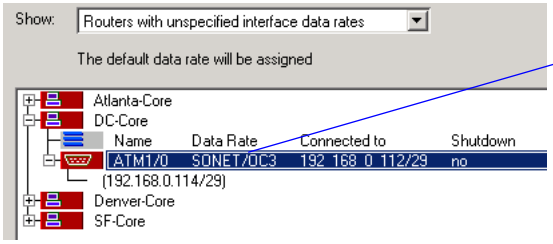

If you do not specify a data rate, the default rate is assigned. The default rate for this ATM interface is SONET/OC3.

4 Choose a data rate from the Data Rate pull-down menu. If the rate you want does not appear on this menu, choose Edit and specify the data rate (in bps). Then click Apply.

➥ The Data Rate fields for both interfaces are updated to reflect the change, and the interfaces now have light green (OK, modified by user) icons.

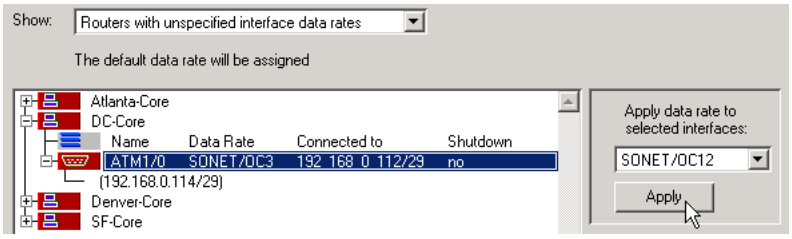

Note—If the icon for an interface is still red after you have specified a data rate, the interface is not connected. Connect the interface as described in Connecting Unnumbered Interfaces on page ISU-10-30, then set the data rate.

End of Procedure 10-6

You can save the changes you made to a model assistant file, which you can then use during another import. To create a model assistant file from the changes made in the Import Assistant, click Save changes after you have specified all of the missing information. The model assistant file is created in the primary models directory with the name you specify. For information about editing model assistant files, see Model Assistant on page ISU-6-31 of the Sentinel *User Guide*.

Viewing the Import Summary Report

When the import is complete, a concise import summary opens. The report contains single-line descriptions and counts of significant events that occurred during the import.

Figure 10-20 Concise Import Summary Report

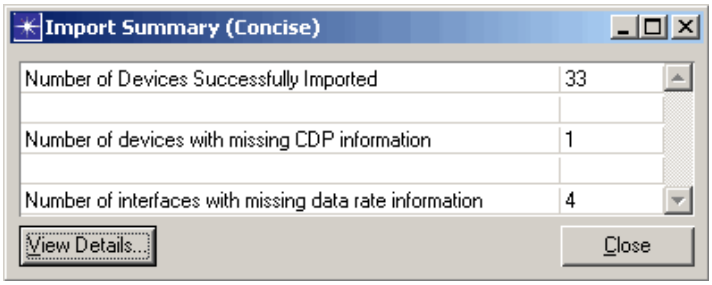

By resting the cursor over each item, you can see expanded information about the event(s) in the form of a tooltip.

The View Details button, when clicked, opens a full import summary report.This report contains more information about the import including the timestamp, duration of process, specified import mode and options selected. It also shows the directory from which the device configuration files were imported and issues that were identified by the Import Assistant during the process.

Figure 10-21 Import Summary Report

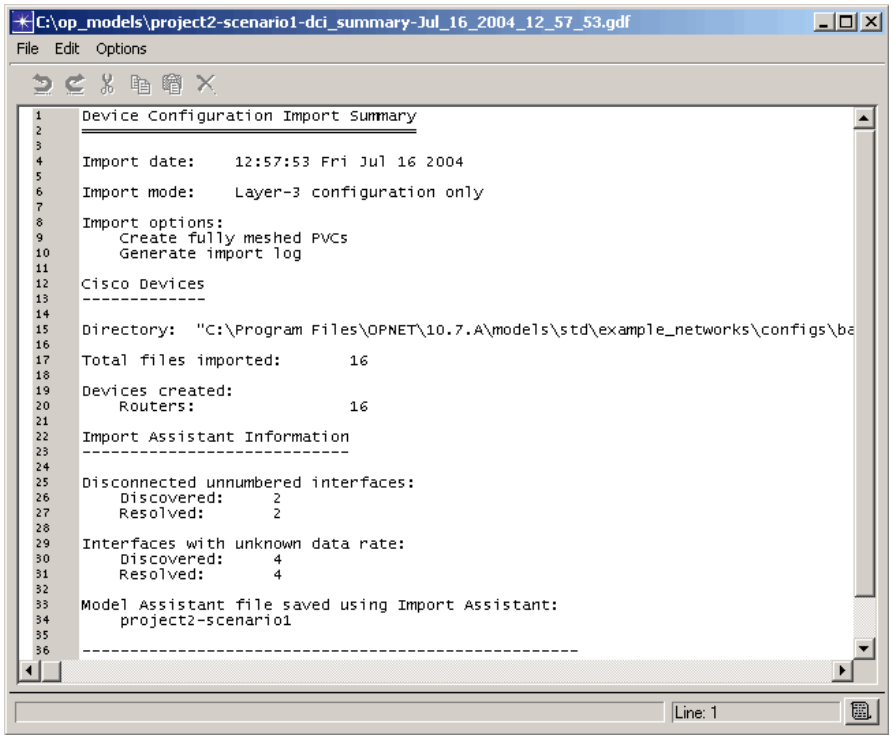

If you wish to review this import summary or other import summaries for previous imports, right-click in the workspace and choose View Import Summary. A dialog box opens that contains a list of import summary reports from previous device configuration imports. You can select any of the available reports for viewing.

Figure 10-22 View Import Summary

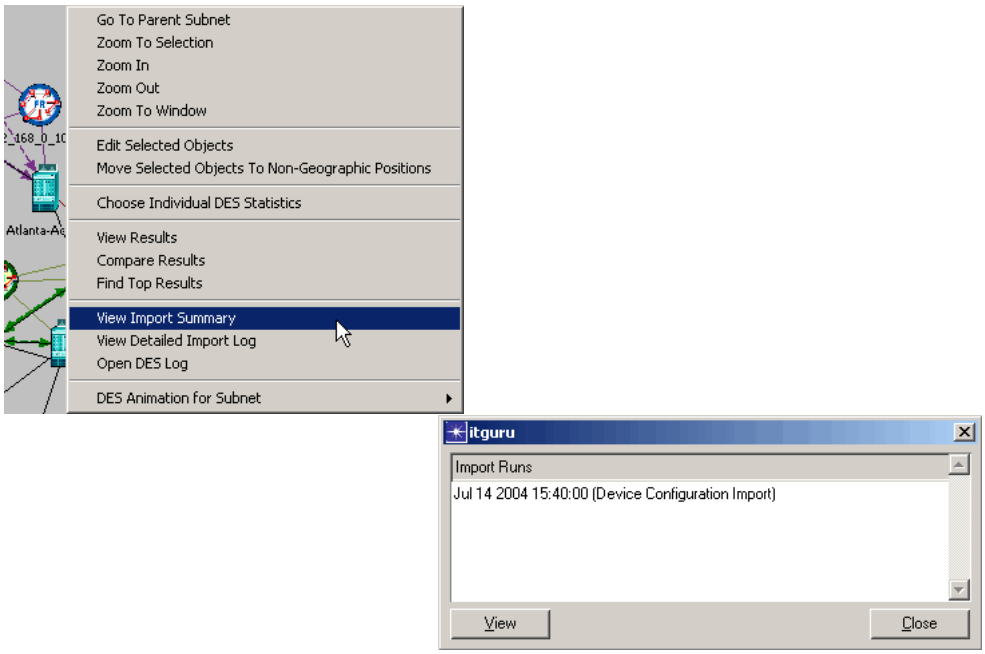

Using Model Assistant Files with Imported Topologies

Model assistant files let you specify information about the network topology that is not included in the configuration files.

To apply model assistant files after importing a network, use the procedure described in Model Assistant on page ISU-6-31 of the Sentinel *User Guide*. You can also apply model assistant files during the import, by specifying the files in the Import Device Configurations dialog box. If you specify information about unnumbered interfaces and unspecified data rates in model assistant files, applying the files during the import lets you bypass the Import Assistant.

For general information about how to create, edit, and apply model assistant files, see Model Assistant on page ISU-6-31 of the Sentinel *User Guide*.

After the Import

Although you have imported a virtual model of your network, the live network continues to evolve. OPNET allows you to manage the network model along with the changes to your “real” network.

- Incremental Imports and Updates to Network Model
- Viewing the Import Log
- View and Edit Configuration Files
- Virtual Command Line Interface
- Network Inventory Summary Report

Incremental Imports and Updates to Network Model

You can use DCI to update an existing network topology, whether or not that topology was created with DCI. You can

- add new devices to the network
- update the configuration files of some devices
- do a new import for the entire network

An import of a network model subsequent to an initial import is also called an *incremental import*. There are several ways to perform an incremental import. Refer to Procedure 10-4 and note the following import modes related to incremental imports:

- Merge with Existing Model—Use this option to add devices to the existing topology. If some of the devices are already in the topology, these devices are replaced (that is, reimported) using the configuration files of this import. DCI matches the devices based on the hostname of the node.
- Reimport Configurations for Selected Devices—Use this option to update the configuration files of a few devices in the existing topology. The devices must already be selected in the workspace. If the devices you want to update are not currently selected, cancel this operation, select the devices, and start over.

You do not need to specify the location of the configuration files; DCI uses the file specified in the Import Information attribute on the device. By default, this is the location specified in the original import.

- Reimport Configuration for Modified Devices—When you select this option, DCI will reimport configuration files for any devices that have been modified since the last import, based on comparison of the timestamp of the last import and the current time.

Note—When DCI updates the configuration files of nodes that were created during a previous import, it deletes the existing node and creates a new node in the network based on the information in the updated configuration file. You will lose any changes you made to the node in the model since the previous import. You can eliminate the need to reconfigure each device manually after an update by using model assistant files, as described in Using Model Assistant Files with Imported Topologies on page ISU-10-34. If you make the original changes in a model assistant file, you can use this file when updating configuration files in later imports.

The import mode, Reimport Configuration for Modified Devices, is the fastest method, especially for large networks. If the network is relatively small, or if you have many changes, consider replacing the entire model. All options use the basic import procedure described in Procedure 10-4. They differ only in the import mode used in the Import Device Configurations dialog box.

Table 10-5 Import Mode Usage for Updating Imported Topologies

To...	Use This Import Mode...
Add new devices to an existing network	Merge with Existing Model
Update the configuration files of some nodes	Reimport Configurations for Selected Devices
Perform incremental import of changes in the network model	Reimport Configuration for Modified Devices
Reimport the files of all devices in the network	Replace Entire Model
End of Table 10-5	

When adding new devices to an existing network or updating the configuration files for a few devices in the network, create a directory that contains only the configuration files of the new devices you want to import. Specify this directory in the Import Device Configurations dialog box and make sure that you are using the import mode specified in Table 10-5.

Viewing the Import Log

During the import, some commands in the device configuration files may have been skipped or ignored. You can see which commands were skipped during the import by selecting View Device Configuration Import Log from the workspace pop-up menu. You can view the import log for an individual node by selecting View Device Configuration Import Log from the node's object pop-up menu. In the import Log Browser for both types of logs (global and node), you can sort the skipped commands according to line number, command (Message column), class, and so on.

Figure 10-23 DCI Import Log

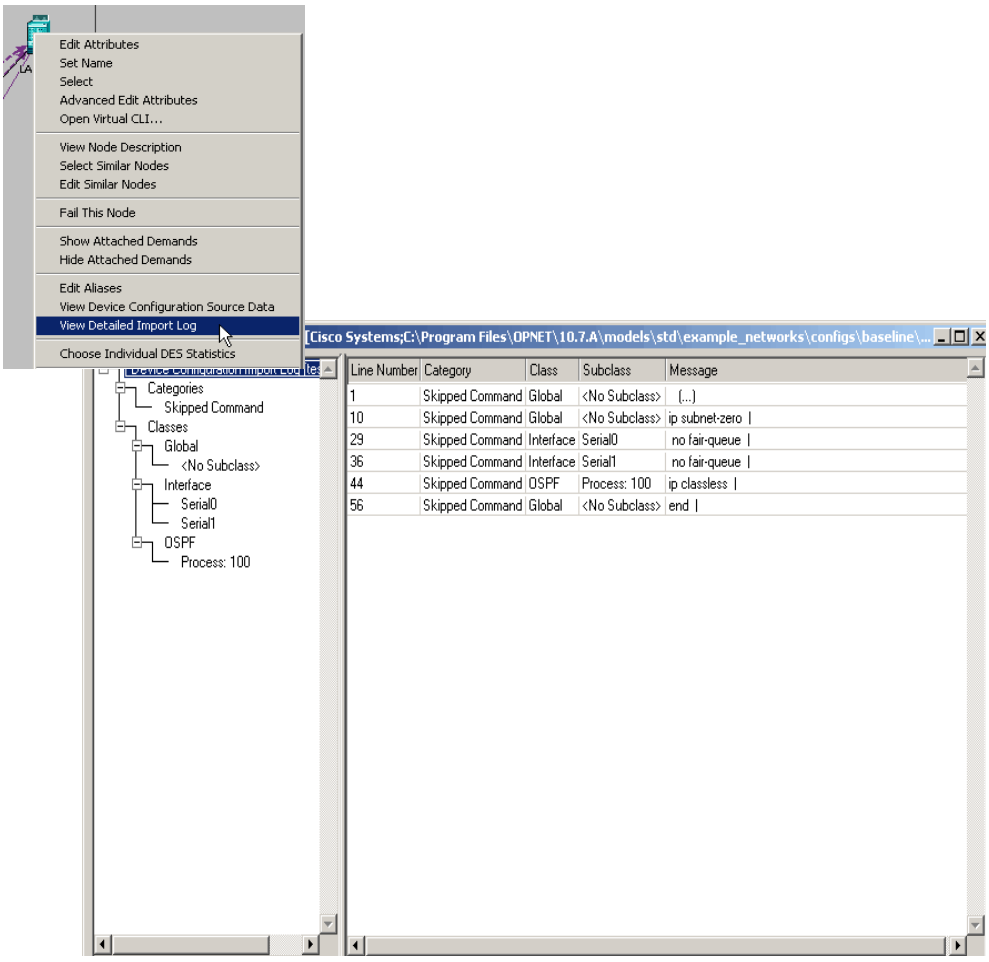

Viewing and Editing Configuration Files After an Import

After an import, you can view the configuration file that DCI used to create a node in the Project Editor. To do this, select View Device Configuration Source Data from the node's object pop-up menu.

Procedure 10-7 Editing a Configuration File of a Device

1. Right-click on the device in the Project Editor workspace and select View Device Configuration Source Data from the pop-up menu.

   ➥ The configuration file appears in a text edit pad (see Figure 10-24 on page ISU-10-38).

2. Edit the configuration file.

3. Choose File > Save to save your changes.

4. Close the text edit pad.

End of Procedure 10-7

After editing a configuration file in IT Sentinel, you need to import the updated configuration file so that the changes you made are included in the topology. See Incremental Imports and Updates to Network Model on page ISU-10-35 for more information.

Figure 10-24 Viewing the Device Configuration File After Import

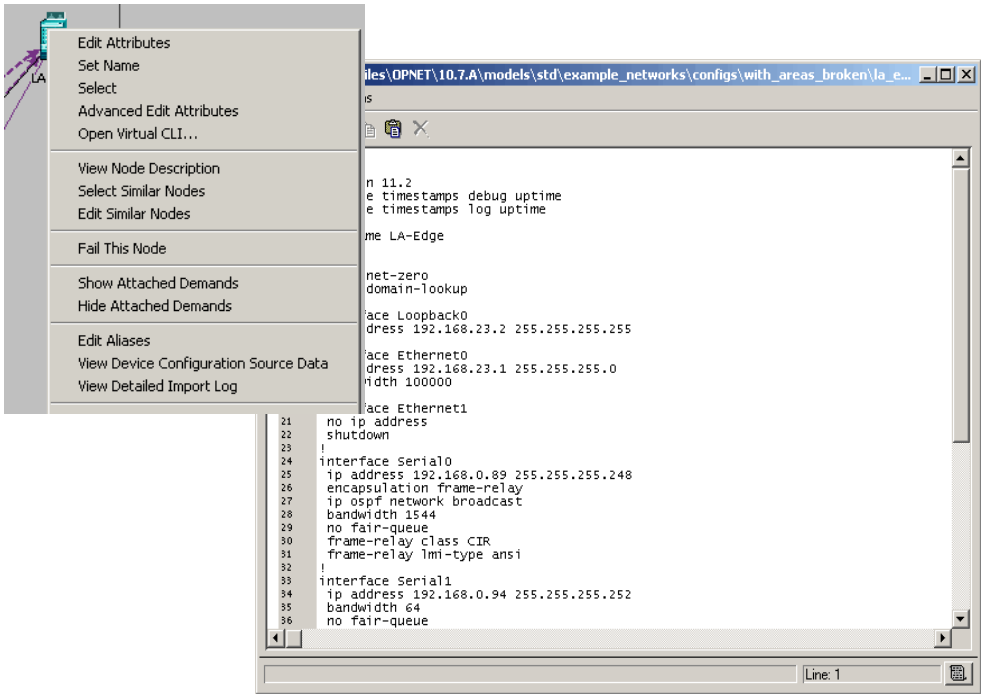

Using the Virtual Command Line Interface

The last section described how to permanently edit your configuration files through a text editor. If you want to edit the devices only for a particular project, you can use Virtual Command Line Interface (Virtual CLI), which allows you to edit Cisco routers and switches within the network model. The changes made with Virtual CLI are specific to the project within which you are working and do not affect the underlying source files.

The Virtual CLI has the same look and feel as the Cisco CLI, although it supports a sub-set of Cisco commands. Virtual CLI supports familiar Cisco CLI usage including shortcuts (such as typing `ena` instead of `enable`), command help access via "?", and tab completion of a command.

To access the Virtual CLI, right-click on a Cisco router or switch and select Open Virtual CLI... Figure 10-6 shows an example of a Virtual CLI session.

Figure 10-25 Virtual Command Line Interface

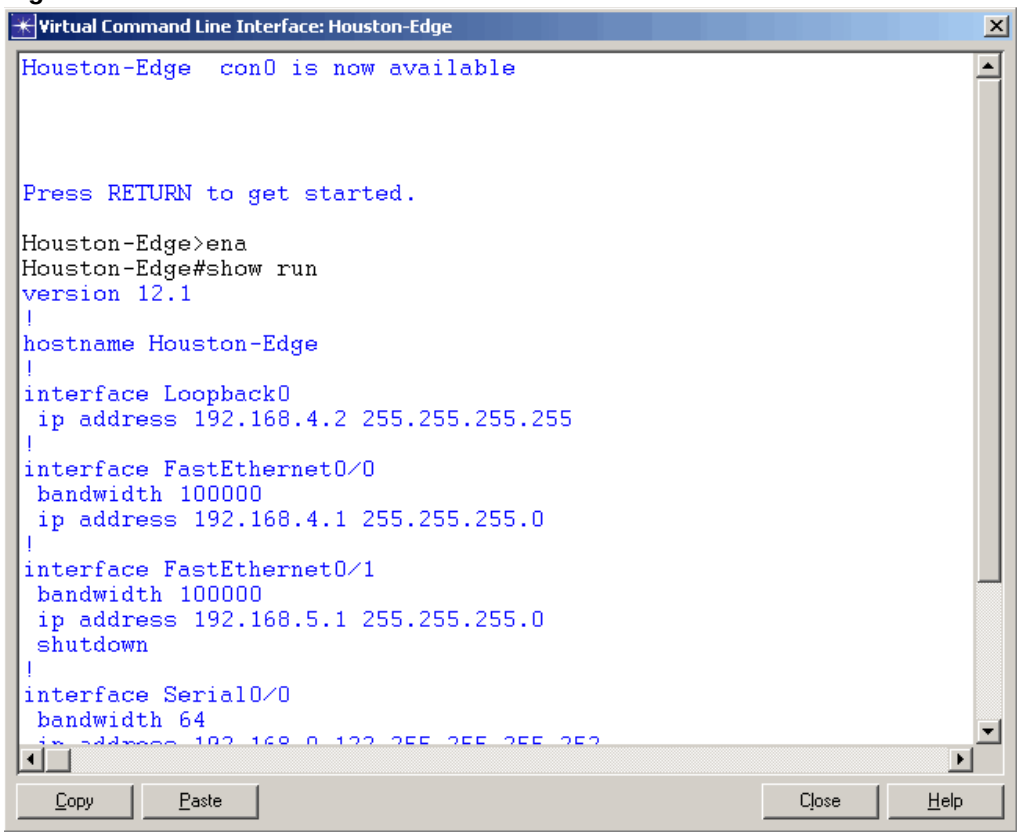

Network Inventory Summary Report

After an import, and on a periodic basis, you may want to see how many devices are in your network and of what type. If you have a rather small imported network model, it is a simple task to visually gather this information. For larger networks, however, it is a time-consuming, error-prone task to take inventory. The Network Inventory Summary report is available at any time after an import. Running a Network Inventory Summary report after an incremental import allows you to keep your inventory files up-to-date. In Figure 10-26 you can see that there are two ways to generate the report: through the drop-down menus or through the toolbar.

Figure 10-26 Generating a Network Inventory Summary

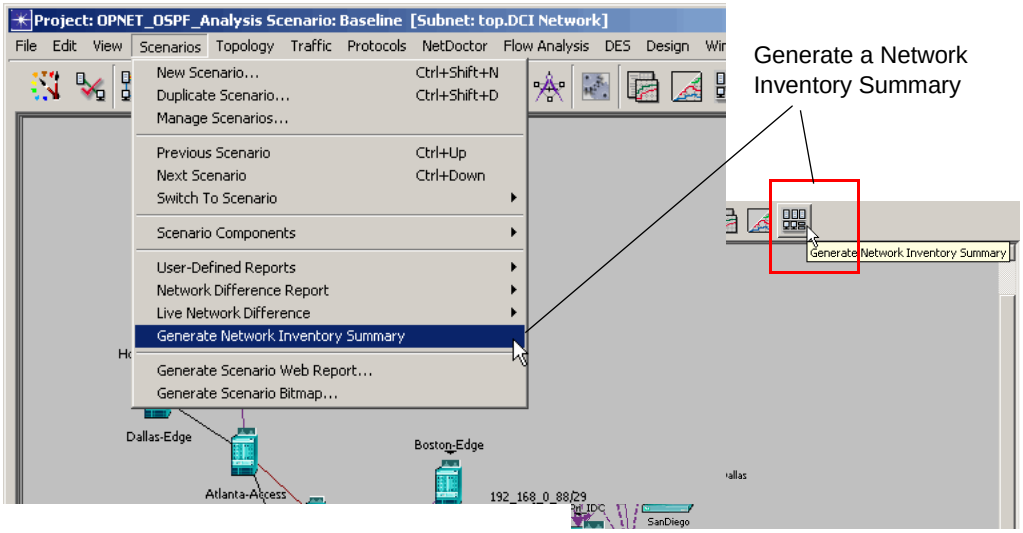

An example of a Network Inventory Summary report is shown in Figure 10-27. The summary includes

- Devices—A total count of devices, and a breakdown by device type
- Vendors—A count by vendor
- Physical Links—A total count of physical links, and a breakdown by type
- Virtual Links—A total count of virtual links, and a breakdown by type
- Other—A total count of other types of devices in the model, including network clouds, application models, configuration utilities, etc.

Figure 10-27 Network Inventory Summary

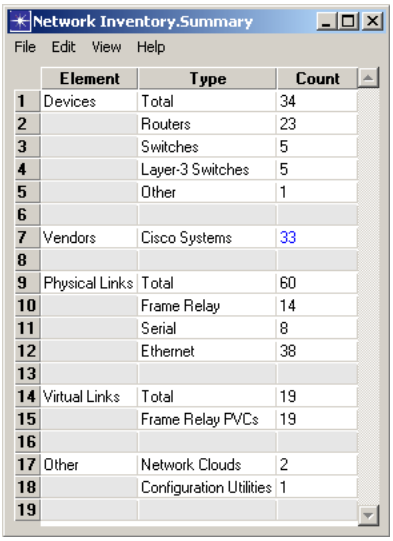

	Element	Type	Count
1	Devices	Total	34
2		Routers	23
3		Switches	5
4		Layer-3 Switches	5
5		Other	1
6			
7	Vendors	Cisco Systems	33
8			
9	Physical Links	Total	60
10		Frame Relay	14
11		Serial	8
12		Ethernet	38
13			
14	Virtual Links	Total	19
15		Frame Relay PVCs	19
16			
17	Other	Network Clouds	2
18		Configuration Utilities	1
19			

If a count listing is in blue text, there are further actions you can take. In the example shown here, you can also drill down into the Cisco listing for specifics about the devices found. Notice in the Device Details dialog box, shown in Figure 10-28, more counts are in blue text.

Figure 10-28 Vendor Detail

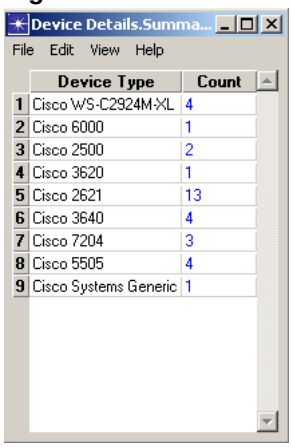

	Device Type	Count
1	Cisco WS-C2924M-XL	4
2	Cisco 6000	1
3	Cisco 2500	2
4	Cisco 3620	1
5	Cisco 2621	13
6	Cisco 3640	4
7	Cisco 7204	3
8	Cisco 5505	4
9	Cisco Systems Generic	1

To drill down even further, you can click on any number in blue. In the example shown in Figure 10-29, the four Cisco WS-C2924M-XL switches are listed.

Figure 10-29 Detail of Cisco Model WS-C2924M-XL

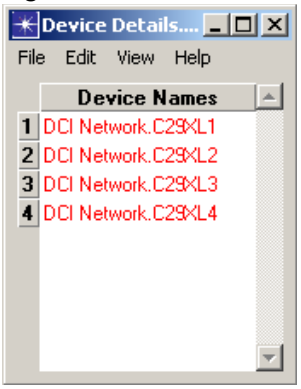

At any point in the summary, you can also export the report to any of the following file types:

- Spreadsheet (for example, MS-Excel)
- Comma-Separated ASCII
- Web Report
- XML

Figure 10-30 Export Network Inventory Summary

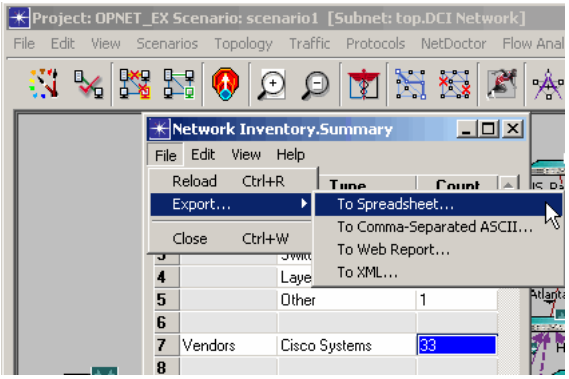

Exporting the report makes it easier to use the information for analysis, management reporting, and other requirements. Figure 10-31 shows an example of this Network Inventory Summary report exported to a web report format.

Figure 10-31 Export to a Web Report

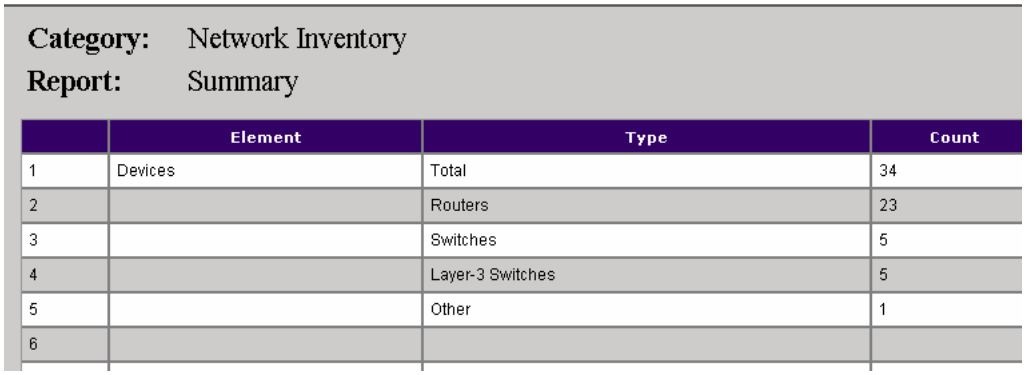

Category: Network Inventory
Report: Summary

	Element	Type	Count
1	Devices	Total	34
2		Routers	23
3		Switches	5
4		Layer-3 Switches	5
5		Other	1
6			

11 Publishing Results

This chapter describes the different reports that IT Sentinel can generate. For more information, see:

- Generate Scenario Web Report
- User-Defined Reports on page ISU-11-3
- Network Difference Reports on page ISU-11-5

Generate Scenario Web Report

To generate a scenario web report manually, or create an automation task step for web reports, choose Scenarios > Generate Scenario Web Report. The resulting web report includes a set of HTML and GIF files that show every subnetwork, object, and attribute list in the scenario.

You can traverse this report in a Web browser in much the same way you traverse a scenario in the Project Editor, as shown in Figure 11-1. Click on each subnet to enter that subnetwork, and click on each link or node object to view its attributes. Use the browser's Back key to return to the previous view.

Figure 11-1 Viewing a Network Topology in a Web Browser

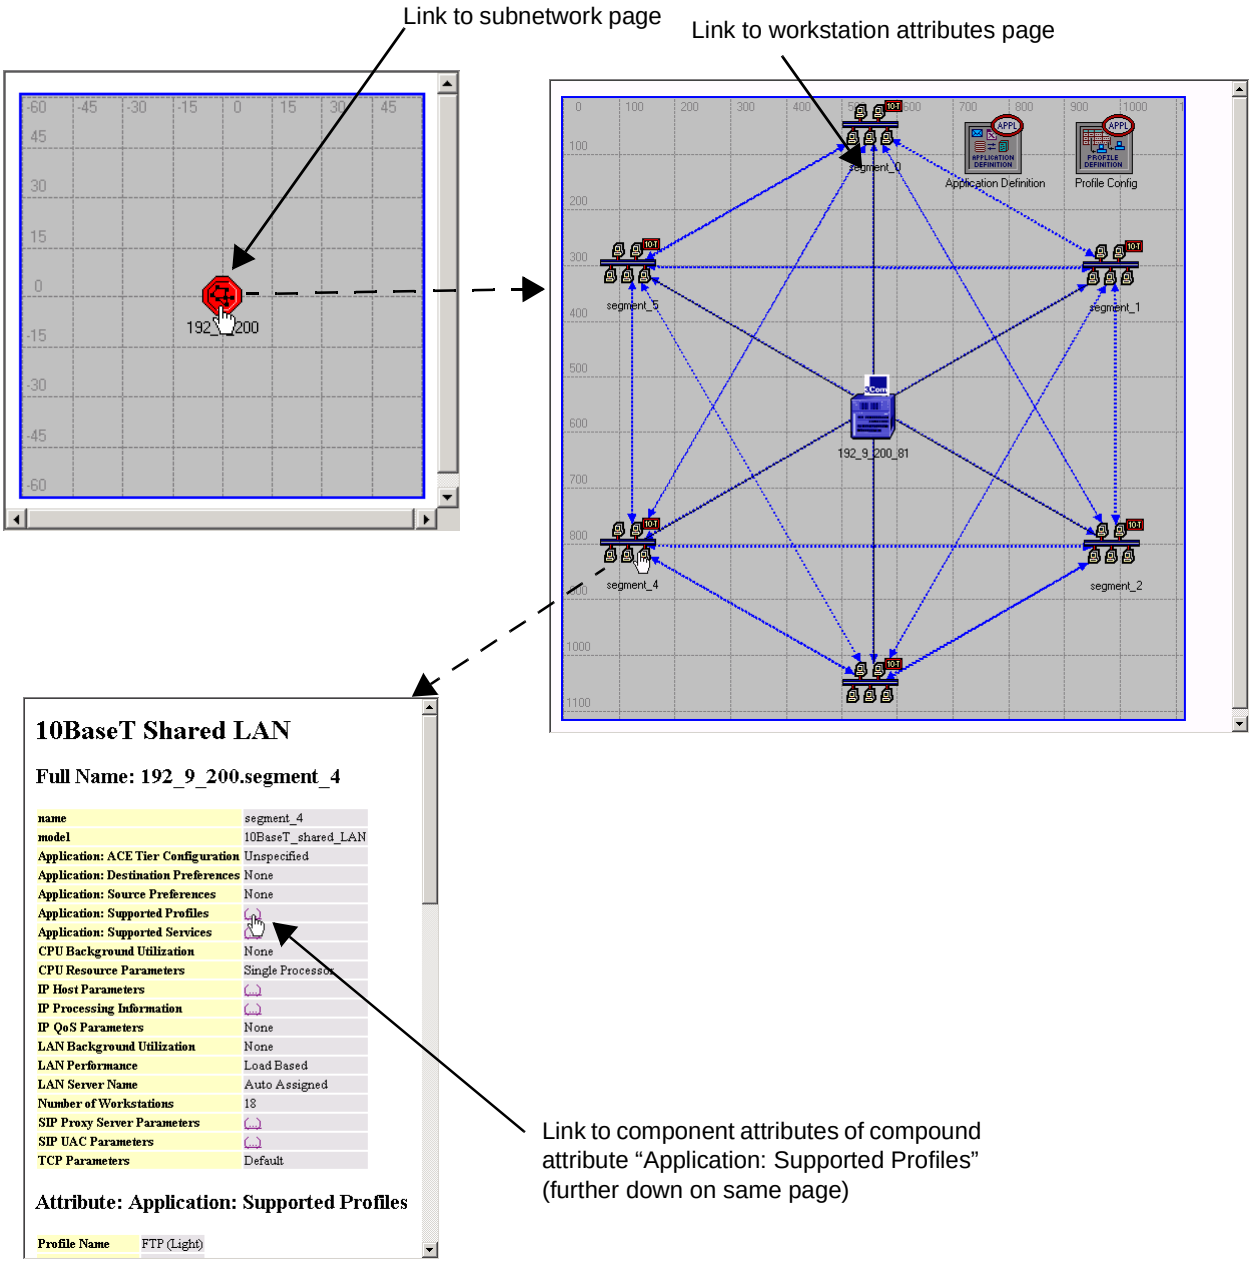

User-Defined Reports

A *user-defined report* is a set of output tables that list objects and attribute values of interest. User-defined reports are useful when you want to view and compare the attributes of multiple objects in custom tables. There are two kinds of user-defined reports:

- Static reports contain a snapshot of attribute values at the time of the report. Several tables can be specified and generated at one time. Static reports can be viewed on-screen or included in a web report.
- Live reports contain continuously updated attribute values and dynamically selected objects.

The two kinds of user-defined report tables have a similar appearance (see Figure 11-2). In both, you can click on an object name (shown in red) to zoom the network workspace to the corresponding object. The main difference is that the contents of a live report will change dynamically as you select different objects in the network or edit attribute values.

Figure 11-2 User-Defined Report

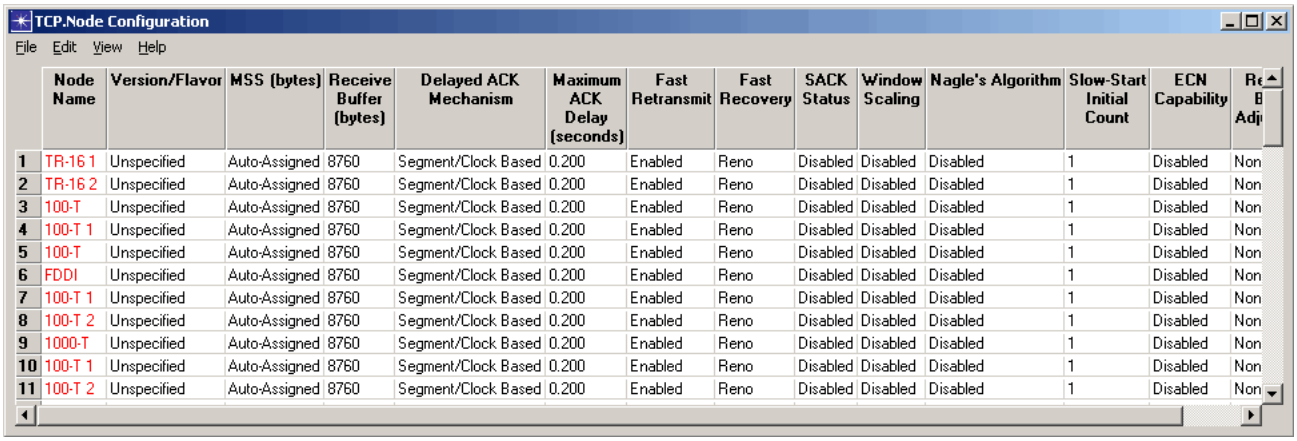

	Node Name	Version/Flavor	MSS (bytes)	Receive Buffer (bytes)	Delayed ACK Mechanism	Maximum ACK Delay (seconds)	Fast Retransmit	Fast Recovery	SACK Status	Window Scaling	Nagle's Algorithm	Slow-Start Initial Count	ECN Capability	R… B… Adj…
1	TR-16 1	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
2	TR-16 2	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
3	100-T	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
4	100-T 1	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
5	100-T	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
6	FDDI	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
7	100-T 1	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
8	100-T 2	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
9	1000-T	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
10	100-T 1	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non
11	100-T 2	Unspecified	Auto-Assigned	8760	Segment/Clock Based	0.200	Enabled	Reno	Disabled	Disabled	Disabled	1	Disabled	Non

You control what kinds of information a report includes by specifying an *attribute template*. You can use predefined attribute templates from the OPNET Model Library or from some other source. You also can define custom attribute templates that include the specific objects and attributes you are interested in (see Defining an Attribute Template on page ISU-9-10).

Follow Procedure 11-1 to generate a static report or Procedure 11-2 to open a live report.

Procedure 11-1 Generating a User-Defined Report

1 If you want to include a custom table in your report, choose Edit > Edit Attribute Template and define the table as described in Defining an Attribute Template on page ISU-9-10.

2 Choose Scenarios > User-Defined Reports > Generate Report from Template.

➥ The Generate User-Defined Report dialog box appears.

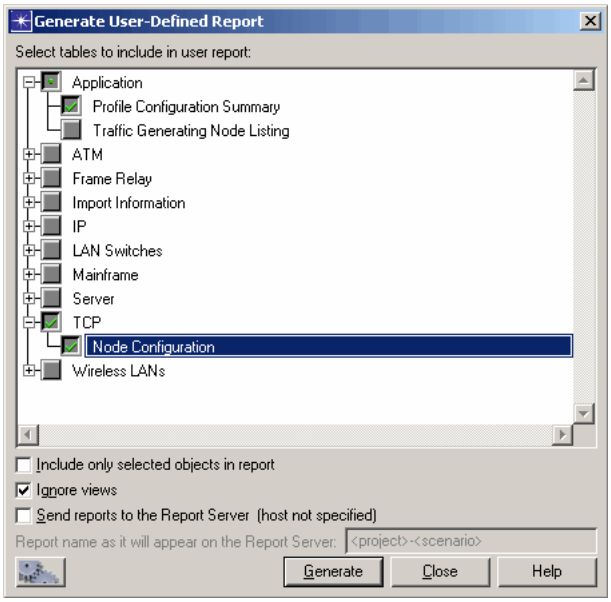

3 In the treeview, select the tables you want to include in the report.

4 Set the following options, as needed:

- Include only selected objects in report—When this check box is selected, only selected objects are included in the report.
- Ignore views—When this check box is selected, the report considers all objects in the scenario, regardless of whether they are included in the current view.
- Send reports to the Report Server (<hostname>)—When this check box is selected, a copy of the generated report is sent to the Report Server specified by the report_server_name preference (Report Server module required).
- Report name as it will appear on the Report Server—Identifying name for the report sent to the Report Server (available only when the Report Server option is selected).

  You can use two variables (<project> and <scenario>) in this field. For example, you could enter “Mary’s <project>-<scenario>”. If you run a report from project “Corporate” and scenario “Internetwork”, the resulting report will be named “Mary’s Corporate-Internetwork”.

5 Click Generate.

➥ The View Results window opens. The treeview lists all tables included in the report.

- To open a table in a separate window, select the table name and click Show. The Preview pane will show a partial version of the selected table.
- To create a web report of all tables, click Generate Web Report.

End of Procedure 11-1

Procedure 11-2 Opening a Live User-Defined Report Table

1 If you want to open a custom table, first choose Edit > Edit Attribute Template to define the table (as described in Defining an Attribute Template on page ISU-9-10).

2 Choose Scenarios > User-Defined Reports > Open Live Report Table.

➥ The Live User-Defined Report Tables dialog box appears.

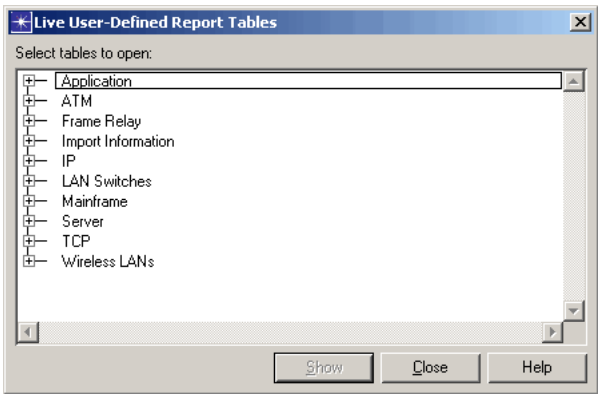

3 In the treeview, select the table you want to open and click Show.

➥ The live report table appears.

Note—If no objects matching the table definition were selected in the network, the table will be empty.

4 Select the network objects you want to view in the report. You can:

- select objects one at a time to see the attributes of each sequentially
- select multiple objects and scroll through the table to compare attributes

Use any of the available selection methods (such as click-and-drag or the Select Similar Nodes operation) to select objects of interest.

End of Procedure 11-2

Network Difference Reports

A *network difference report* compares objects and attributes between two scenarios. Network difference reports are useful when you want to see exactly how a network has evolved from one scenario to the next or how networks imported from different sources vary. A difference report lists

- All objects that appear in both networks but have different attribute settings
- All objects that only appear in the current network (the scenario from which the report was generated)

- All objects that only appear in the other network
- All objects that have identical attributes in both networks

Note—IT Sentinel determines the presence of a network object by comparing the full hierarchical names of objects in both scenarios. The full hierarchical name includes the object name and every parent subnet. An object is considered to be the “same object” in both scenarios if and only if the object has the same hierarchical name in both scenarios.

Procedure 11-3 Generating a Network Difference Report

1. Open one of the scenarios you want to compare. Normally you want to open the “after” scenario and compare it to the “before” or baseline scenario.
2. To define the objects and attributes you want to exclude from the report, choose Scenarios > Network Difference Report > Define Report. (For more information, see Defining a Difference Report on page ISU-11-6.)
3. To generate the difference report, open the Generate Difference Report dialog box. To open this dialog box, do one of the following:
   - Choose Scenarios > Network Difference Report > Generate Report
   - In the Define Difference Report dialog box, click the Generate button.
4. Specify the second scenario to compare, the difference behavior file, and the target directory for the difference data (for more information, see Table 11-1 Define Difference Report Dialog Box on page ISU-11-7). Then click Generate

   ➥ IT Sentinel compares the two scenarios and generates one or more XML data files to store the results.

   ➥ The Network Difference Tree Viewer shows the results of the difference operation. For more information, see Viewing a Difference Report on page ISU-11-10.

End of Procedure 11-3

Defining a Difference Report

You can exclude some types of objects and attributes from a comparison report. For example, you might want to ignore differences in annotations or wireless domains. Similarly, you might want to ignore specific types of attributes. The resulting report is more concise and easier to navigate.

To exclude objects or attributes from a difference report, choose Scenarios > Network Difference Report > Define Report. This opens the Define Difference Report dialog box, shown in Figure 11-3 on page ISU-11-7.

Define Difference Report Dialog Box

Figure 11-3 Define Difference Report Dialog Box

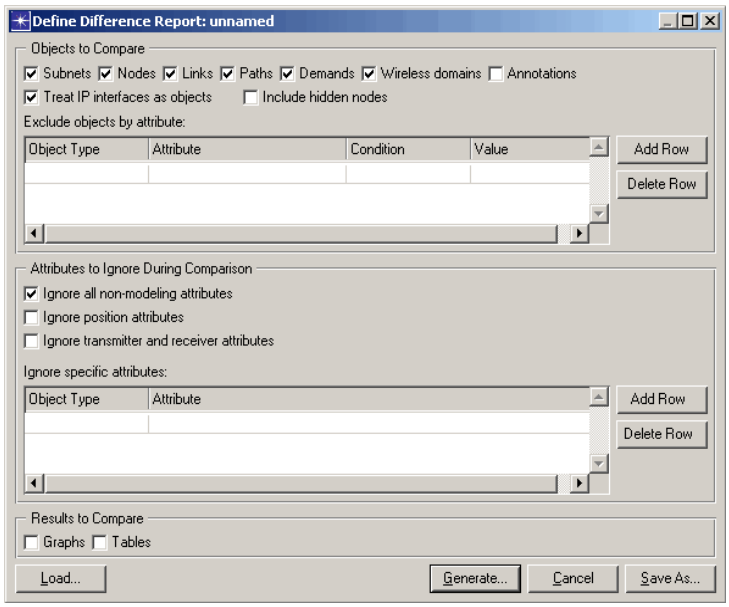

Table 11-1 Define Difference Report Dialog Box (Part 1 of 2)

Option	Description
Attributes to ignore during comparison	Specify attributes to ignore: • All non-modeling attributes—Ignore attributes that do not affect discrete event simulation or Flow Analysis results (examples: “color”, “line style”, “creation source”) • Position attributes—Ignore attributes that specify physical locations of objects (examples: “x position”, “y position”, “x span”, “y span”) • Transmitter and receiver attributes—Ignore link attributes that describe how the link attaches to its end-point objects. (examples: “transmitter”, “transmitter a”, “receiver b”)
“Ignore specific attributes” table	Ignore specific attributes during the comparison. Attributes are specified by a hierarchical path (example: IP Routing Parameters ⧫ OS Version).
Results to Compare	Specify whether to include output results in the report: • Graphs—Compare generated graphs in each scenario and report differences, if any. • Tables—Compare generated tables in each scenario and report differences, if any.
Save As	Save the current settings to a behavior file
Load	Load report settings from a behavior file

Table 11-1 Define Difference Report Dialog Box (Part 2 of 2)

Option	Description
Generate	Generate a difference report based on the current settings
End of Table 11-1	
Attributes to ignore during comparison	Specify attributes to ignore: • All non-modeling attributes—Ignore attributes that do not affect discrete event simulation or Flow Analysis results (examples: “color”, “line style”, “creation source”) • Position attributes—Ignore attributes that specify physical locations of objects (examples: “x position”, “y position”, “x span”, “y span”) • Transmitter and receiver attributes—Ignore link attributes that describe how the link attaches to its end-point objects. (examples: “transmitter”, “transmitter a”, “receiver b”)
End of Table 11-1	

Generating a Difference Report

To generate a network difference report, choose Scenarios > Network Difference Report > Generate Report. You can also generate a report by clicking the Generate button in the Define Difference Report dialog box. Either action opens the Generate Difference Report dialog box, shown in Figure 11-4.

Figure 11-4 Generate Difference Report Dialog Box

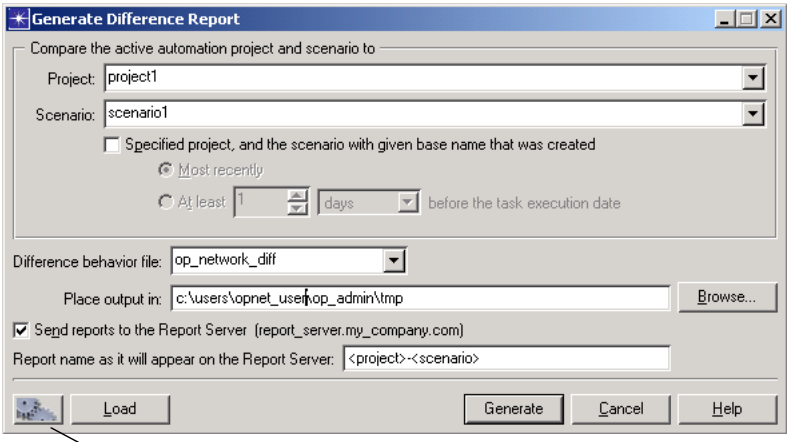

Table 11-2 Generate Difference Report Dialog Box

Option	Description
Compare the active project and scenario to...	Choose the project and scenario with which you want to compare the active scenario.
Specified project, and ...	Compare the active scenario with a scenario created before the task execution date
Difference behavior file	Use the difference criteria defined in the selected behavior file. You can create a difference behavior file in the Define Difference Report dialog box; for more information, see Defining a Difference Report on page ISU-11-6
Place output in	Target directory for the generated difference data.
Send reports to...	Send a copy of the generated report to the Report Server (specified by the report_server_name preference)
Report name...	File name of the report sent to the Report Server. You can use two variables (“<project>” and “<scenario>”) in this field. Thus, you could enter the following string: “Mary’s <project>/<scenario>” . If you run a report from project Corporate and scenario Internetwork, the resulting report will be named “Mary’s Corporate/Internetwork”. This option is available only when the “Send reports to...” option is selected.
Save Settings for Automation	Save the current settings to an automation (.af) file
Load	Load settings from an automation (.af) file
Generate	This button issues the command to generate the report.
End of Table 11-2	

Viewing a Difference Report

The Network Difference Tree Viewer (Figure 11-5) opens automatically after you generate a difference report. You can also view the last generated report by choosing Scenarios > Network Difference Report > View Last Report.

Figure 11-5 Network Difference Tree Viewer

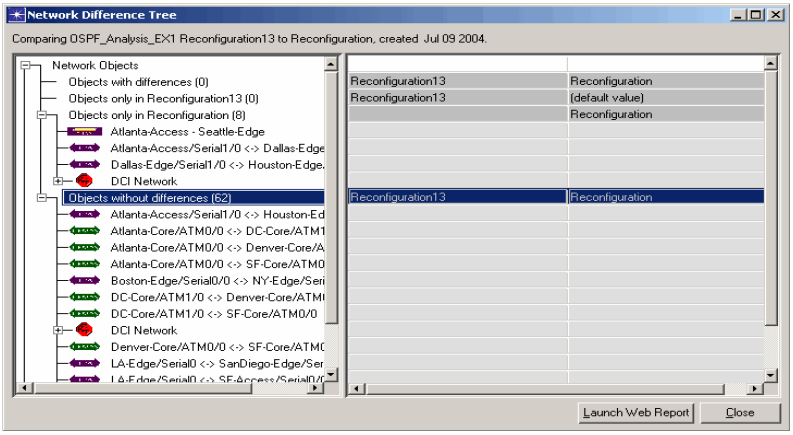

You can generate a web report from this viewer; IT Sentinel stores the results in the same directory as the difference data files. To view the last generated web report, choose Scenarios > Network Difference Report > View Last Web Report.

Index

Symbols

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

R

S

T

U

V

W

Z