NETWARS

Routing and Transport Protocols

2005-1 Final User's Guide (OPNET 236)

Contract DASW01 03 D 0008

October 17, 2005

Prepared for:
Defense Contracting Command -
Washington
Washington, DC 50310-5200

Prepared by:
OPNET Technologies, Inc.
7255 Woodmont Avenue
Bethesda, MD 20814-7904

Identification

Document Identification

Document Title: Routing and Transport Protocols
Version: Final

Software Identification

Product Name: NETWARS
Product Release: 5.1

Documentation Conventions

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

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

Objects, Examples, and System I/O

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

```
C:\Netwars\User_Data\Projects
```

- 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 (< >):

```
<NETWARS path>\Scenario_Builder\op_admin\err_log
```

Object Hierarchies

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

Edit > Preferences > Advanced

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.

Document Revision History

Release Date	Product Version	Chapter	Description of Change
October 17, 2005	5.1 Final	Front material	Changed refs in Documentation Conventions section to be NETWARS-specific rather than OPNET-specific (i.e, changed `<opnet_user_home>\`... ref to `<NETWARS path>\Scenario_Builder\`...)
		1	Added Using Virtual CLI section.
		2	In Procedure 2-1, updated Topology > Import > From Device Configuration Files to Topology > Import > Device Configuration Files.
August 3, 2005	5.1 Draft	All	Preliminary version.

Contents

List of Figures

List of Procedures

1 Introduction

NETWARS Overview

The Command, Control, Communications, and Computer Systems Directorate of the Joint Staff, in partnership with the Defense Information Systems Agency, Directorate for Technical Integration Services, developed Network Warfare Simulation (NETWARS). NETWARS provides modeling and simulation (M&S) capabilities for measuring and assessing information flow through strategic, operational, and tactical military communications networks. Analyzing the results from NETWARS can provide considerable utility in determining which communication systems might be overloaded during selected times in a particular scenario, and can assist with making prudent acquisition planning decisions.

Document Overview

This user's guide, *Routing and Transport Protocols,* is a technology tutorial, not a tutorial on the use of NETWARS. It provides hands-on exercises that illustrate the basic functioning of the following protocols:

- OSPF and IS-IS,
- EIGRP and RIP, and
- IP Multicast Networks.

The examples provided in this guide were presented at OPNETWORK as part of various routing and transport protocol workshops. If you do not have access to the sample files referenced in a particular example, simply view the screenshots provided in this guide.

Referenced Documents

- *Standard Model User Guides*, OPNET Technologies.

Routing Protocols

OPNET provides behavioral models for various routing protocols. These models provide excellent support in discrete event simulation (DES) and traffic flow analysis for studying:

- Routing Tables,
- Convergence Statistics,
- Effects of configuring areas in a network,
- Failure/Recovery Analysis, and
- Traffic Engineering Techniques.

Chapter 2 discusses the use of OSPF (Open Shortest Path First) and IS-IS (Intermediate System-to-Intermediate System), which are link state (LS) routing protocols.

Chapter 3 discusses the use of EIGRP (Enhanced Interior Gateway Routing Protocol) and RIP (Routing Information Protocol), which are distance vector-based routing protocols.

Link State Routing Protocols

OSPF and IS-IS are very similar routing protocols. Each node maintains a view of the entire network topology, Route Table computation is centralized, and convergence is fast due to the fact that every node knows the entire topology.

But OSPF and IS-IS have some differences. OSPF is standardized by IETF (RFC 2328), more widely deployed in enterprises, and only used for routing IP. IS-IS is standardized by ISO (ISO 10859), more widely deployed in ISPs, and used for routing IP and CLNP (also adapted for IPX.)

Chapter 2 provides example network scenarios using OSPF and IS-IS.

Distance Vector Routing Protocols

With distance vector routing protocols such as EIGRP and RIP, each node maintains distance to each destination, Route Table computation is distributed, and convergence can be slow due to the fact that changes in topology can take longer to propagate.

Chapter 3 provides example network scenarios using EIGRP and RIP.

IP Multicast Networks

OPNET also models multicast routing protocols. Multicasting is a method for delivering content from a single sender to multiple receivers. It’s different from broadcasting, where traffic is sent to all possible receivers (sending email to everyone in a company). Multicasting sends traffic only to receivers who have shown interest in receiving traffic destined for a specific multicast address.

The benefits of using multicasting include:

- Reduce network load,
- Conserve network bandwidth,
- Easier information sharing infrastructure, and
- Send to a single group address and be done.

OPNET support for multicast includes DCI support for various Cisco and Juniper commands, Netdoctor multicast rule suite, and Simulation (DES) support.

Chapter 4 provides example scenarios using a multicast network.

Using Virtual CLI

All Cisco router node models have a feature that allows you to configure the router model using a replica of the Cisco IOS command line interface. The Virtual Command Line Interface (or 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**. You can then enter commands as you would if you were working on a router console.

Procedure 4-6 in Chapter 4 provides an example of using Virtual CLI to fix configuration errors.

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.

2 Planning and Analyzing OSPF and IS-IS Networks

Introduction

OSPF and IS-IS are common Interior Gateway Protocols (IGPs) used for routing in IP networks. OSPF and IS-IS are link state routing protocols developed by IETF and ISO respectively. Both have fairly complicated algorithms to ensure stability in the network and facilitate faster convergence, and loop-free routing.

In this chapter, we will:

- Study the behavior of OSPF in a flat network (no areas),
  - — Observe equal cost multi-path (ECMP) routing,
  - — Study routing tables and LSDB size,
  - — Analyze network convergence and effects of timers, and
  - — Perform basic traffic engineering.
- Study the behavior of OSPF in a network with areas, and
- Merge IS-IS areas by adding a NET.

We will conclude that:

- OSPF and IS-IS route by exchanging link state information,
- You can perform basic traffic engineering by modifying costs (auto-calculate costs based on bandwidth),
- Convergence time is affected by many factors (you can tweak Hello, Dead and SPF timers to obtain optimal performance),
- Configuring areas and summaries can improve scalability, and
- IS-IS routing is similar to that of OSPF, and merging areas in IS-IS is easy.

Note—The following examples were presented at OPNETWORK 2004 in Session 1313, Planning and Analyzing OSPF and IS-IS Networks. If you do not have access to the files that these procedures use, you can still follow along using the sample screens provided in this user’s guide.

There is a set of models with all of the procedural steps completed. You may use these for reference, if you do not wish to do the exercises yourself. These files end in the suffix “_ref” (reference).

Analysis of Baseline OSPF Network (No Areas)

In this section, we study an OSPF network without areas configured. We measure its performance in regards to routing table sizes and convergence times. At the same time, we examine OSPF features such as multi-path routing.

Procedure 2-1 Create the Virtual Network Model

1 Launch NETWARS, if not already opened.

2 From the System Editor's **File** menu, choose **Open Editor**.

3 From the Open Editor drop-down menu, choose **Scenario Builder**, and then click **OK**.

4 Select **File > Open Project**. The Open Project dialog box displays.

5 In the Open Project dialog box, select the project file named `1313_Planning_and_Analysis_of_OSPF_and_ISIS`, and then click **Open**.

Note—If you do not have access to this file, simply view the screens provided in this user's guide to follow along with the procedure.

6 Build the network topology from router configuration files.

6.1 Choose **Topology > Import > From Device Configuration Files**.

6.2 Select **Replace entire model**; checkmark **Cisco Router IOS**, and type in the path to the `configs_baseline` directory, as shown below.

6.3 Under Model Assistant Files, choose `1313_Planning_and_Analysis_of_OSPF_and_ISIS-Lab1-ma_export`.

6.4 Click **Import**.

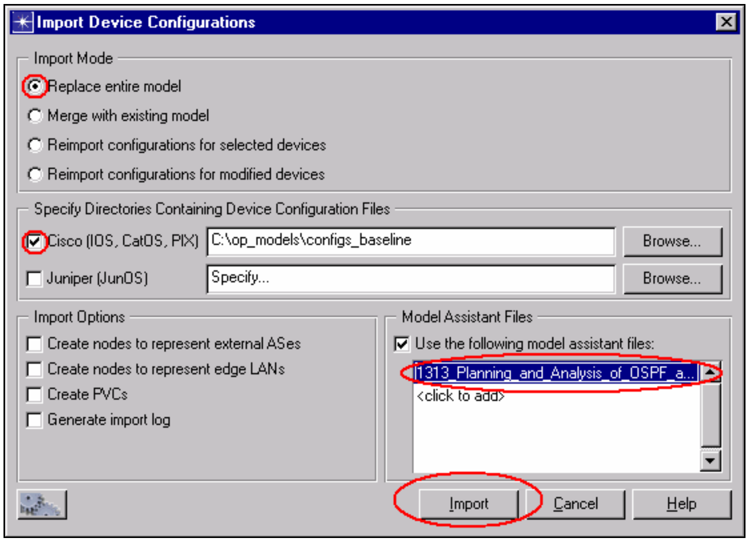

Import Device Configurations

Import Mode

● Replace entire model

○ Merge with existing model

○ Reimport configurations for selected devices

○ Reimport configurations for modified devices

Specify Directories Containing Device Configuration Files

☑ Cisco (IOS, CatOS, PIX) C:\op_models\configs_baseline Browse...

☐ Juniper (JunOS) Specify... Browse...

Import Options

☐ Create nodes to represent external ASes

☐ Create nodes to represent edge LANs

☐ Create PVCs

☐ Generate import log

Model Assistant Files

☑ Use the following model assistant files:

1313_Planning_and_Analysis_of_OSPF_a...

<click to add>

Figure 2-1 Import Device Configurations dialog box

6.5 After the import completes, click **OK** to close the Import Summary window. The following network is created from the given router configuration files.

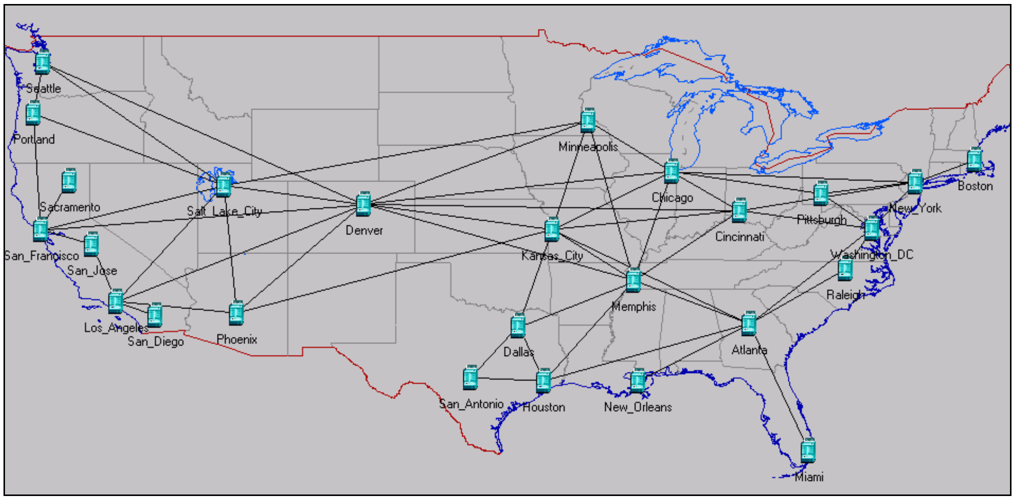

Figure 2-2 Imported OSPF Network Model

This scenario contains 26 routers, each configured with OSPF as their routing protocol. This is a flat network, with no configured areas. All routers are in the backbone area (area 0).

End of Procedure 2-1

Procedure 2-2 Configure the Imported OSPF Network

1 Import traffic flows into the network.

1.1 Traffic data has been provided in a spreadsheet format for import; choose **Traffic > Import Traffic Flows > From Spreadsheet ...**

1.2 Navigate to your NETWARS `op_models` directory.

1.3 Select the file `1313_Lab1a_traffic.txt`; click **Open**.

1.4 Leave the **Replace all existing traffic** radio button checked; click **OK**.

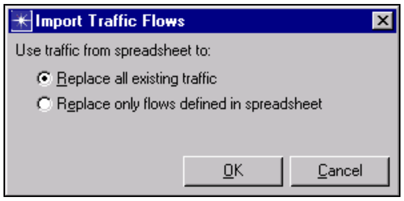

Import Traffic Flows

Use traffic from spreadsheet to:

● Replace all existing traffic

○ Replace only flows defined in spreadsheet

Figure 2-3 Import Traffic Flows dialog box

1.5 Traffic is imported and a summary log of the import process is provided; click **Close**.

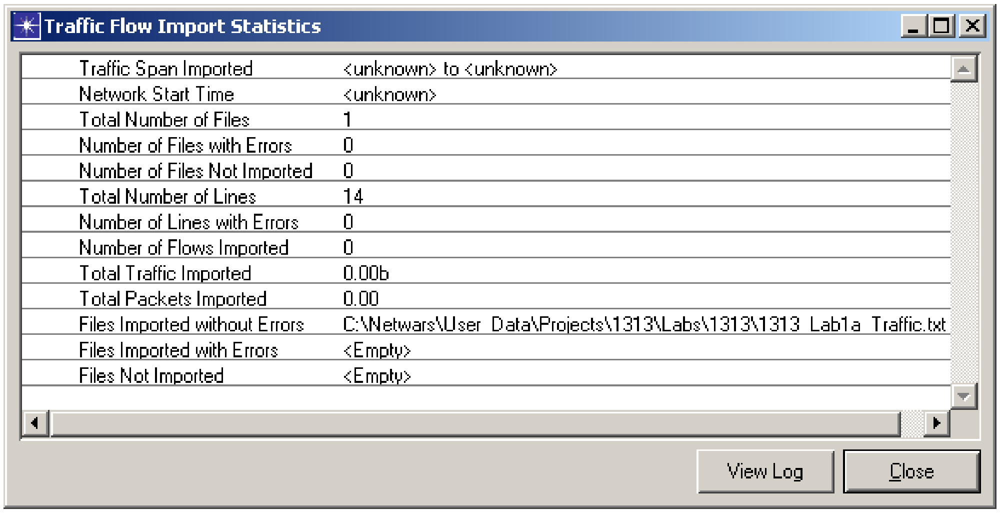

Traffic Span Imported	<unknown> to <unknown>
Network Start Time	<unknown>
Total Number of Files	1
Number of Files with Errors	0
Number of Files Not Imported	0
Total Number of Lines	14
Number of Lines with Errors	0
Number of Flows Imported	0
Total Traffic Imported	0.00b
Total Packets Imported	0.00
Files Imported without Errors	C:\Netwars\User_Data\Projects\1313\Labs\1313\1313_Lab1a_Traffic.txt
Files Imported with Errors	<Empty>
Files Not Imported	<Empty>

Figure 2-4 Traffic Flow Import Statistics dialog box

1.6 Choose **Protocols > IP > Demands > Characterize Traffic Demands…**

1.7 Click the **Record Routes for All Packets** button, and then click **OK**.

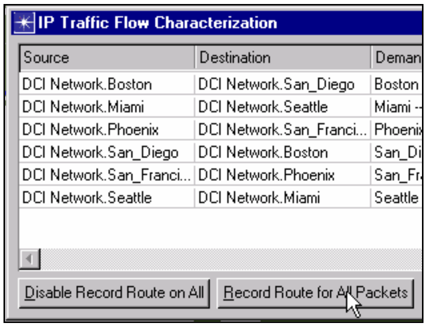
IP Traffic Flow Characterization
Disable Record Route on All
Record Route for All Packets

Source	Destination	Deman
DCI Network.Boston	DCI Network.San_Diego	Boston
DCI Network.Miami	DCI Network.Seattle	Miami -
DCI Network.Phoenix	DCI Network.San_Franci...	Phoenix
DCI Network.San_Diego	DCI Network.Boston	San_Di
DCI Network.San_Franci...	DCI Network.Phoenix	San_Fr
DCI Network.Seattle	DCI Network.Miami	Seattle

Figure 2-5 IP Traffic Flow Characterization dialog box

2 Configure failures in the network.

2.1 Click the **Open Object Palette** button.

2.2 At the top of the Object Palette there is a drop-down list allowing you to switch between available palettes. Click and select the **utilities** palette from the drop-down list.

2.3 Drag-and-drop the Failure Recovery object onto the project space.

2.4 Close the Object Palette.

2.5 Double-click the new Failure Recovery object, then right-click on the Failure Recovery icon and choose **Edit Attributes**.

2.6 Double-click on the Value for Node Failure/Recovery Specification.

2.7 Set the number of rows to **2**.

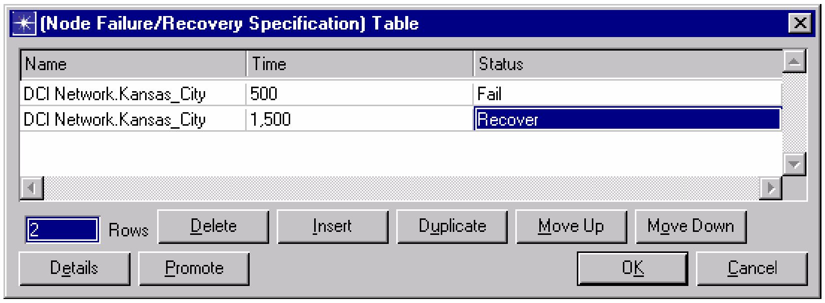

Name	Time	Status
DCI Network.Kansas_City	500	Fail
DCI Network.Kansas_City	1,500	Recover

Rows: 2

Figure 2-6 Node Failure/Recovery Specification Table

2.8 In the first row, choose **Kansas_City** for Name, **500** for Time, and **Fail** for Status; in the second row again choose **Kansas_City** for Name, **1500** for Time, and **Recover** for Status.

2.9 Click **OK** to return to the previous window.

2.10 Double-click on the Value for Link Failure/Recovery Specification.

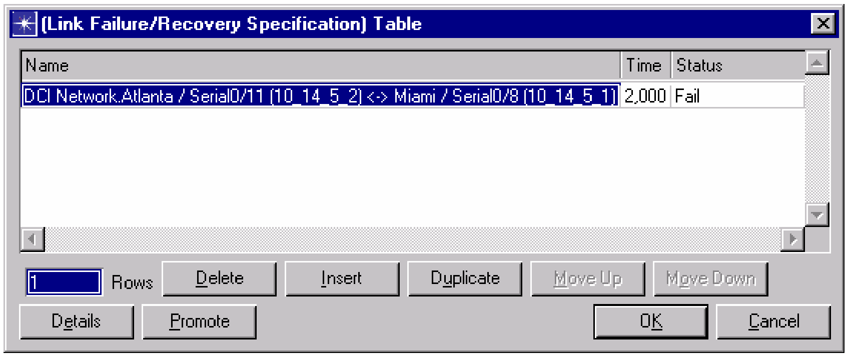

Name	Time	Status
DCI Network.Atlanta / Serial0/11 (10_14_5_2) <-> Miami / Serial0/8 (10_14_5_1)	2,000	Fail

Rows: 1

Figure 2-7 Link Failure/Recovery Specification Table

2.11 Select **1** for Rows, select the name of the link between Atlanta and Miami in the Name field, type **2000** in the Time field, and set Status to **Fail**.

2.12 Click **OK** twice to preserve the changes; as a result of this configuration, the **Kansas_City** router will fail at 500 seconds and recover at 1500 seconds, and the **Atlanta↔Miami** link will fail at 2000 seconds (and will not recover).

3 Configure Load Balancing.

3.1 Choose **Protocols > IP > Routing > Configure Load Balancing Options…**.

3.2 Choose **Packet-Based** for Load balancing options.

3.3 Make sure the **All routers** radio button is selected, and then click **OK**. This will configure routers to perform load balancing for every packet, as opposed to load balancing for every destination.

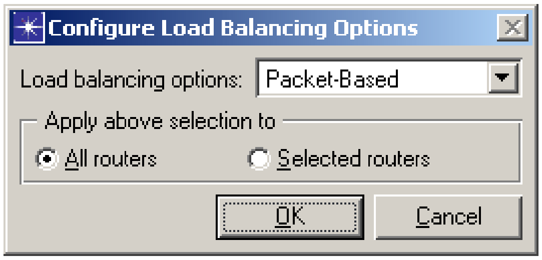

Load balancing options: Packet-Based

Apply above selection to

● All routers ○ Selected routers

Figure 2-8 Configure Load Balancing Options dialog box

4 Configure start time for OSPF on all routers.

In a real network, the OSPF process on different routers would typically start at different times. However for the purpose of this study, we want the OSPF process on all routers in the network to start at the same time. This will eliminate the effects of varying start times on the network convergence time and help us determine precise numbers for the convergence statistics.

4.1 Choose **Protocols > OSPF > Configure Start Time...**.

4.2 In order to start all OSPF processes simultaneously at 5 seconds, choose **constant** for Distribution name, and type **5** for Mean Outcome. Make sure the **All routers** radio button is selected; click **OK**.

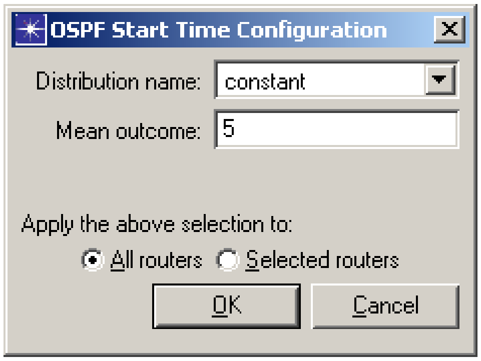

Distribution name: constant

Mean outcome: 5

Apply the above selection to:

● All routers ○ Selected routers

Figure 2-9 OSPF Start TIme Configuration dialog box

End of Procedure 2-2

Procedure 2-3 Configure Collection of Results

1 Export OSPF Link State Database and IP Forwarding Table for all routers.

1.1 Press **Ctrl+A** to select all objects in the network.

1.2 Right-click on the **Seattle** router, and choose **Edit Attributes**.

1.3 Check the **Apply changes to selected objects** checkbox in the lower left corner of the Edit Attributes window.

1.4 Expand the Reports category by clicking on the adjacent (+) button. In the row corresponding to OSPF Link State Database, select **Export at End of Simulation**.

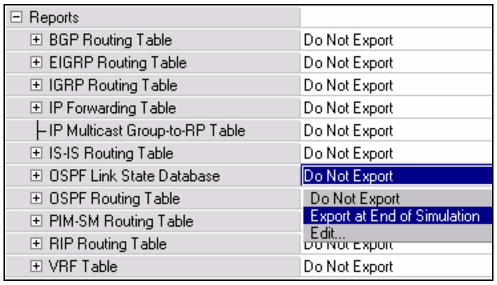

⊟ Reports	
⊞ BGP Routing Table	Do Not Export
⊞ EIGRP Routing Table	Do Not Export
⊞ IGRP Routing Table	Do Not Export
⊞ IP Forwarding Table	Do Not Export
IP Multicast Group-to-RP Table	Do Not Export
⊞ IS-IS Routing Table	Do Not Export
⊞ OSPF Link State Database	Do Not Export
⊞ OSPF Routing Table	
⊞ PIM-SM Routing Table	
⊞ RIP Routing Table	Do Not Export
⊞ VRF Table	Do Not Export

Figure 2-10 Export OSPF Link State Database

1.5 Similar to the previous step, in the row corresponding to IP Forwarding Table, select **Export at End of Simulation**.

1.6 Click **OK**.

1.7 If you get a warning stating that changes to multiple objects cannot be undone, click **Yes**.

1.8 Click somewhere in the project space to de-select all objects.

2 Collect OSPF convergence, traffic and Route Table statistics.

2.1 Right-click in the project space away from all objects, and select **Choose Individual DES Statistics**.

2.2 Expand the Global Statistics category by clicking on the adjacent (+) button.

2.3 Similarly, expand OSPF and select all 3 statistics – Network Convergence Activity, Network Convergence Duration, and Traffic Sent.

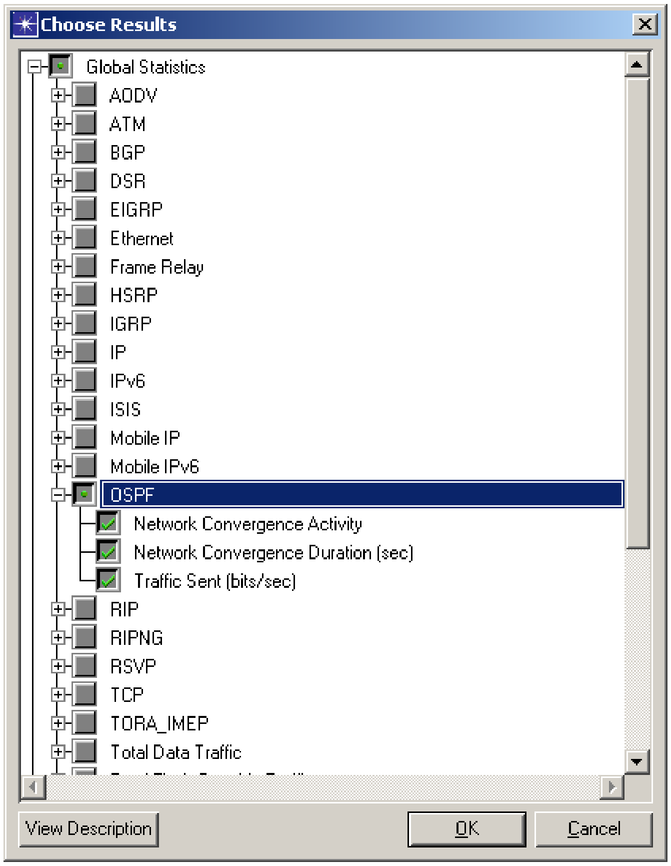

Figure 2-11 Selecting OSPF Statistics

2.4 Expand the Node Statistics category by clicking on the (+) button; expand the Route Table category by clicking on the (+) button.

2.5 Select the Size (number of entries) statistic.

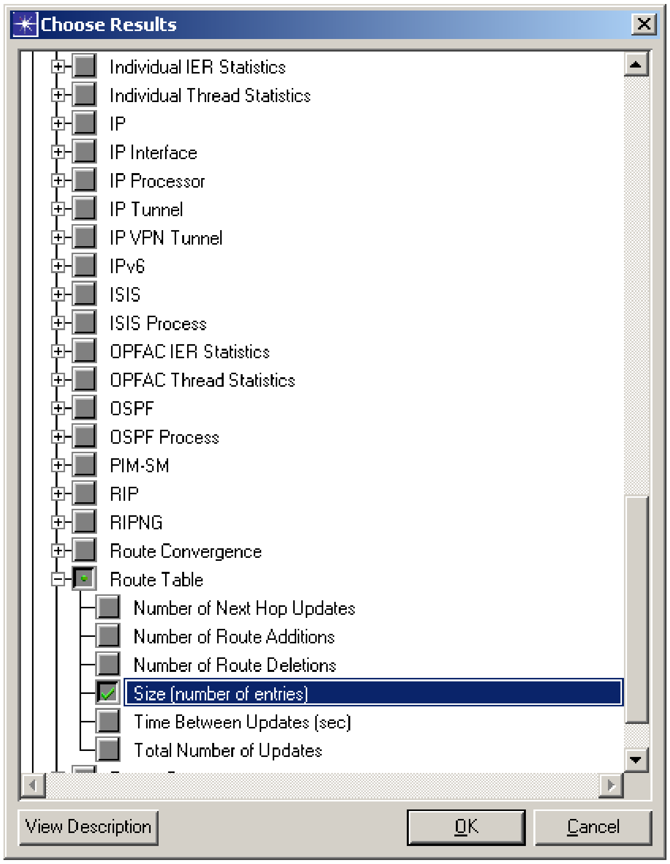

Figure 2-12 Selecting the Route Table Size Statistic

2.6 Click **OK**.

End of Procedure 2-3

Procedure 2-4 Run Simulation and View Results

1 Run the simulation.

1.1 Click the **Configure/Run Simulation** toolbar button.

1.2 In the Inputs section, select the Global Attributes category and scroll down to modify the value of the Tracer Packets Per Interval from **2** to **20**.

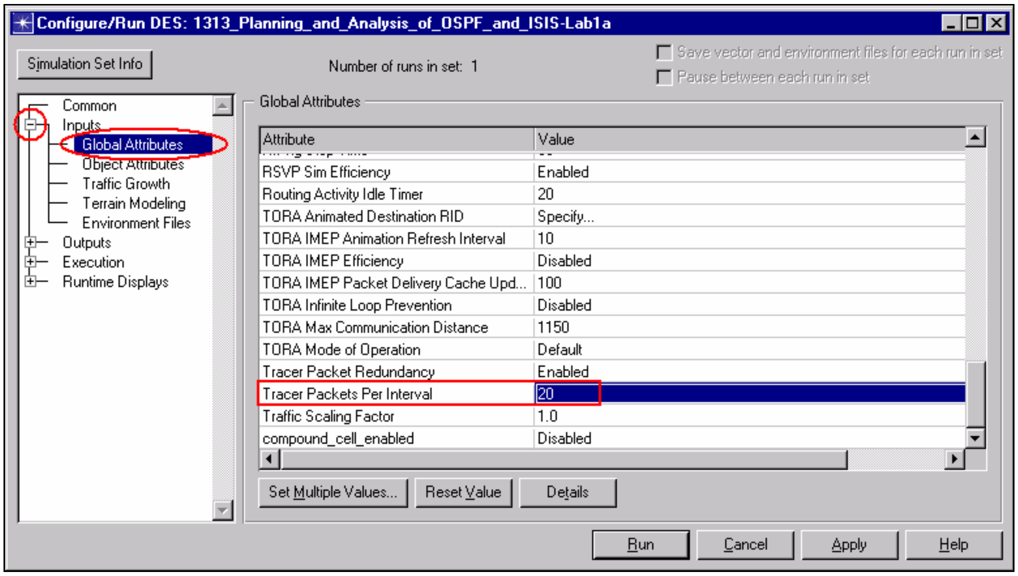

Attribute	Value
RSVP Sim Efficiency	Enabled
Routing Activity Idle Timer	20
TORA Animated Destination RID	Specify...
TORA IMEP Animation Refresh Interval	10
TORA IMEP Efficiency	Disabled
TORA IMEP Packet Delivery Cache Upd...	100
TORA Infinite Loop Prevention	Disabled
TORA Max Communication Distance	1150
TORA Mode of Operation	Default
Tracer Packet Redundancy	Enabled
Tracer Packets Per Interval	20
Traffic Scaling Factor	1.0
compound_cell_enabled	Disabled

Figure 2-13 Configure/Run Simulation dialog box

1.3 In the same Global Attributes category, scroll to the attribute OSPF Sim Efficiency and set it to **Disabled**. Simulation efficiency means that OSPF control traffic will not be simulated after the time specified in the OSPF Stop Time attribute. This is a useful feature for speeding up the simulation, if you know that no further routing changes will occur in the network. Since this is not the case in this exercise, we need to disable OSPF simulation efficiency.

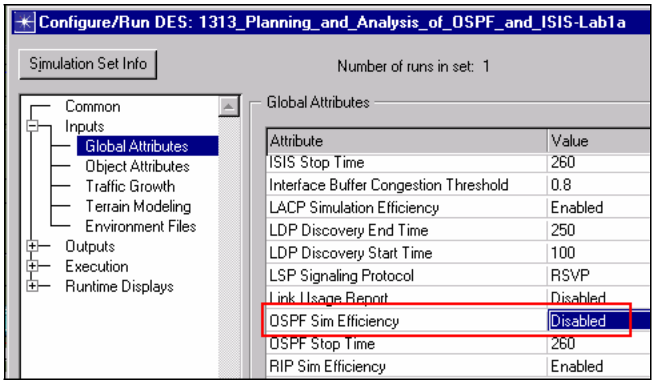

Attribute	Value
ISIS Stop Time	260
Interface Buffer Congestion Threshold	0.8
LACP Simulation Efficiency	Enabled
LDP Discovery End Time	250
LDP Discovery Start Time	100
LSP Signaling Protocol	RSVP
Link Usage Report	Disabled
OSPF Sim Efficiency	Disabled
OSPF Stop Time	260
RIP Sim Efficiency	Enabled

Figure 2-14 OSPF Sim Efficiency Attribute

1.4 Click **Run** to start the simulation.

1.5 Close the Simulation Sequence dialog box after the simulation runs.

2 Inspect the OSPF Link State Database.

2.1 Right-click on the Denver node and select **View Results**.

2.2 Choose the Discrete Event Tables tab.

2.3 Expand the Performance category by clicking on the adjacent (+) button.

2.4 Select OSPF Link State Database Summary at 3600 seconds.

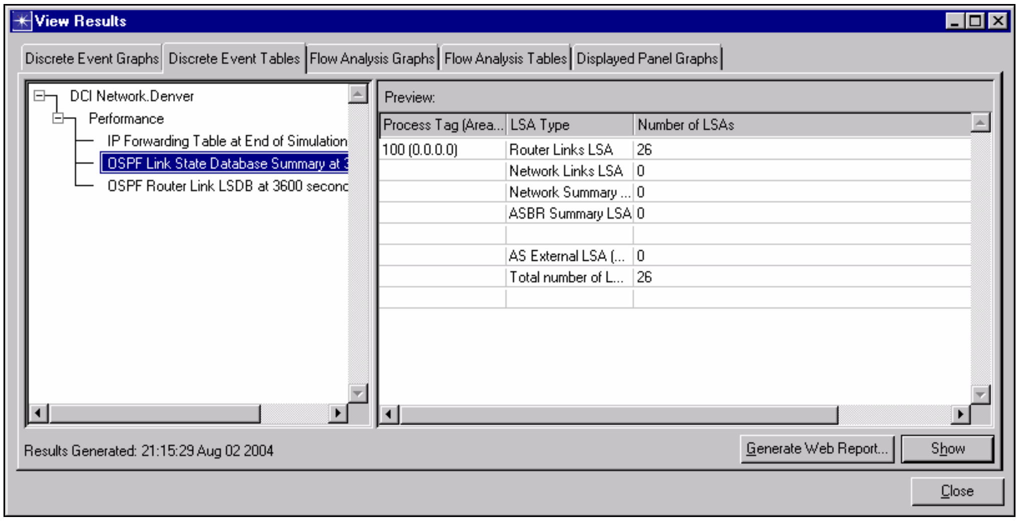

DCI Network.Denver
Performance
IP Forwarding Table at End of Simulation
OSPF Link State Database Summary at 3
OSPF Router Link LSDB at 3600 seconc

Preview:

Process Tag (Area...	LSA Type	Number of LSAs
100 (0.0.0.0)	Router Links LSA	26
	Network Links LSA	0
	Network Summary ...	0
	ASBR Summary LSA	0
	AS External LSA (...	0
	Total number of L...	26

Results Generated: 21:15:29 Aug 02 2004

Figure 2-15 Denver Node’s OSPF Link State Database

2.5 Note that the OSPF Link State Database on **Denver** has LSAs for only one area – Area 0.0.0.0. It has 26 Router Links LSAs and no LSAs of any other type.

2.6 Close the View Results dialog box.

2.7 Now click on the **San_Diego** node and view its OSPF Link State Database summary.

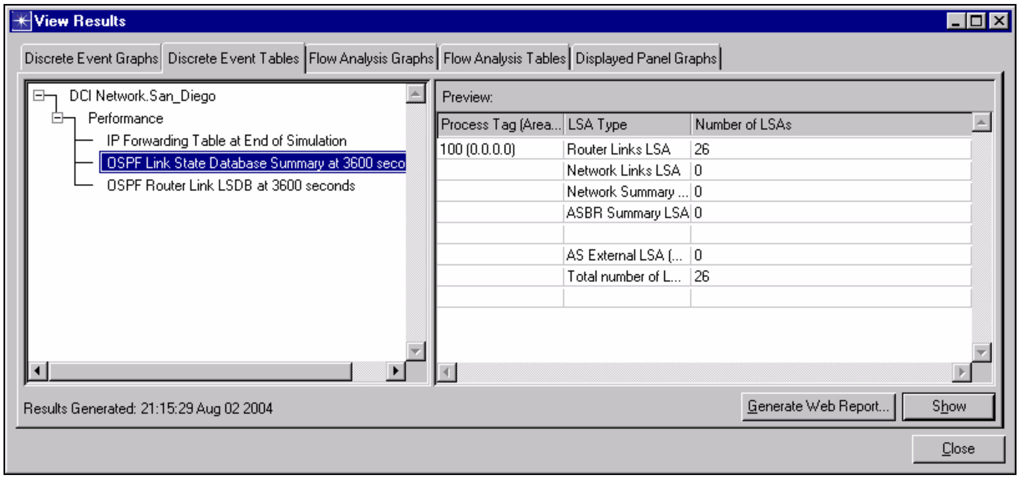

DCI Network.San_Diego
Performance
IP Forwarding Table at End of Simulation
OSPF Link State Database Summary at 3600 seco
OSPF Router Link LSDB at 3600 seconds

Preview:

Process Tag (Area...	LSA Type	Number of LSAs
100 (0.0.0.0)	Router Links LSA	26
	Network Links LSA	0
	Network Summary ...	0
	ASBR Summary LSA	0
	AS External LSA (...	0
	Total number of L...	26

Results Generated: 21:15:29 Aug 02 2004

Figure 2-16 San_Diego Node’s OSPF Link State Database

2.8 Note that the contents of **San_Diego’s** OSPF Link State Database are exactly the same as those of **Denver’s** Link State Database. In fact the Link State Database is exactly the same on all routers in the network.

2.9 Close the View Results dialog box.

3 Examine traffic flow routes at different times.

3.1 Choose **Protocols > IP > Demands > Display Routes for Configured Demands...**.

3.2 Expand the Seattle source by clicking on the (+) buttons, until the **Seattle→Miami** flow is shown. Click on the traffic flow name.

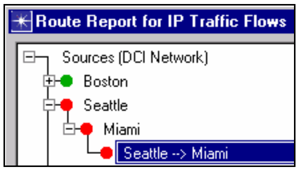

Route Report for IP Traffic Flows

Sources (DCI Network)
Boston
Seattle
Miami
Seattle --> Miami

Figure 2-17 Route Report for IP Traffic Flows dialog box

3.3 On the right side of the panel, each row corresponds to the route over which the traffic flow was routed at a given time. To visualize the routes in the network, click the Display column in a row corresponding to the time of interest to change the field from **No** to **Yes**. Look at the network to visualize the route.

For instance, to display the route at 168.91 seconds, click in the Display field: its value will toggle from **No** to **Yes** and the route will be displayed in the network.

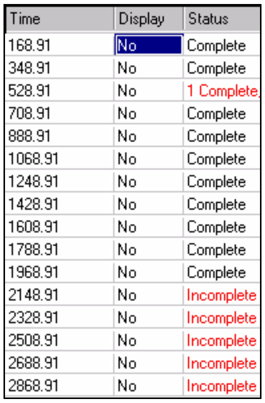

Time	Display	Status
168.91	No	Complete
348.91	No	Complete
528.91	No	1 Complete
708.91	No	Complete
888.91	No	Complete
1068.91	No	Complete
1248.91	No	Complete
1428.91	No	Complete
1608.91	No	Complete
1788.91	No	Complete
1968.91	No	Complete
2148.91	No	Incomplete
2328.91	No	Incomplete
2508.91	No	Incomplete
2688.91	No	Incomplete
2868.91	No	Incomplete

Figure 2-18 Specifying Routes to Display

3.4 Click on the corresponding Display field to display the route taken by the **Seattle→Miami** traffic flow at a time before 500 seconds (when the first failure occurs). Notice that two equal-cost routes are taken by the traffic flow (one through **Kansas_City**, the other through **Memphis**).

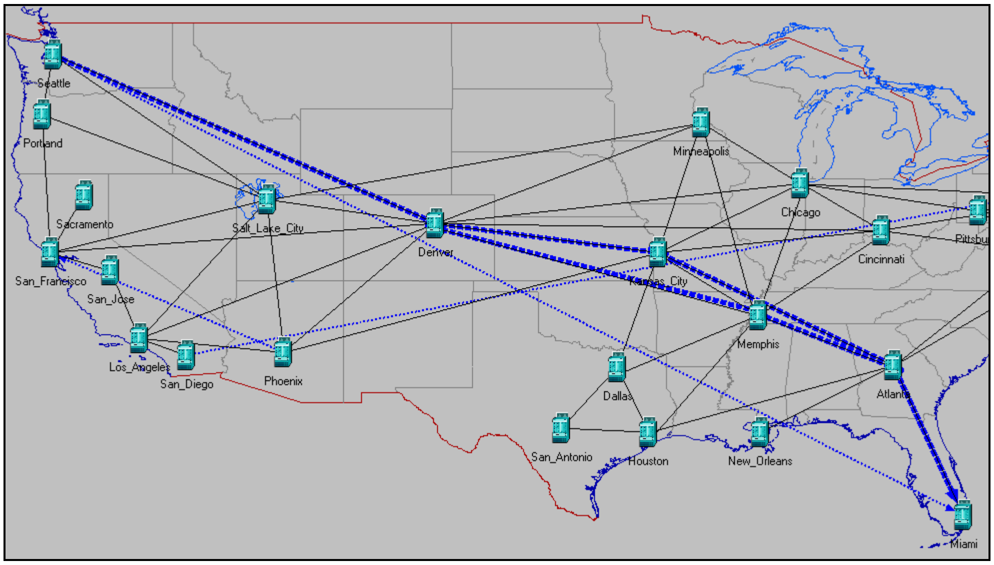

Figure 2-19 Seattle→Miami Route (Before 500 Seconds) Displayed

3.5 Click again on the Display field for the selected time to hide the route.

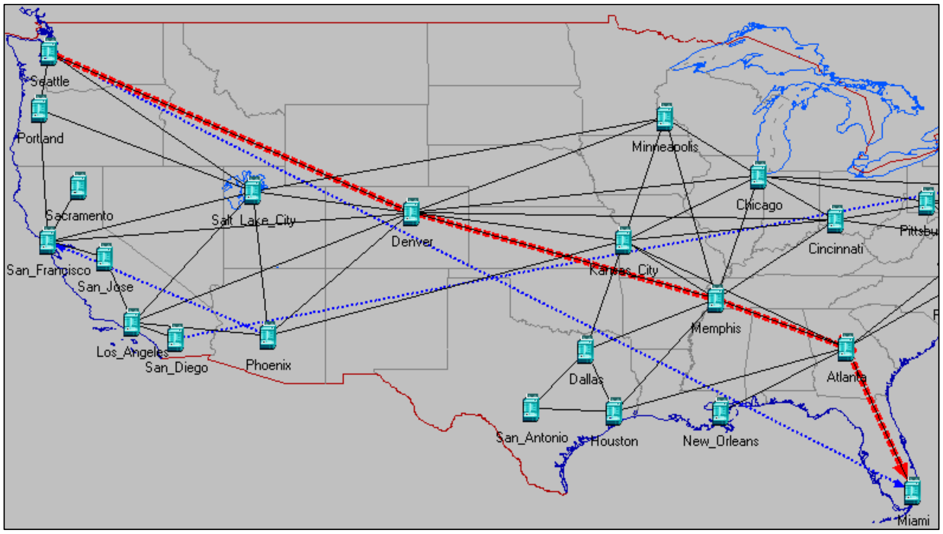

Figure 2-20 Seattle→Miami Route (Before 500 Seconds) Hidden

3.6 Click on the corresponding Display field to display the routes taken by the **Seattle→Miami** traffic flow at a time between 500 and 1500 seconds. Notice that now only one route is available (via **Memphis**) – the route via **Kansas_City** is no longer available since **Kansas_City** node failed at 500 seconds.

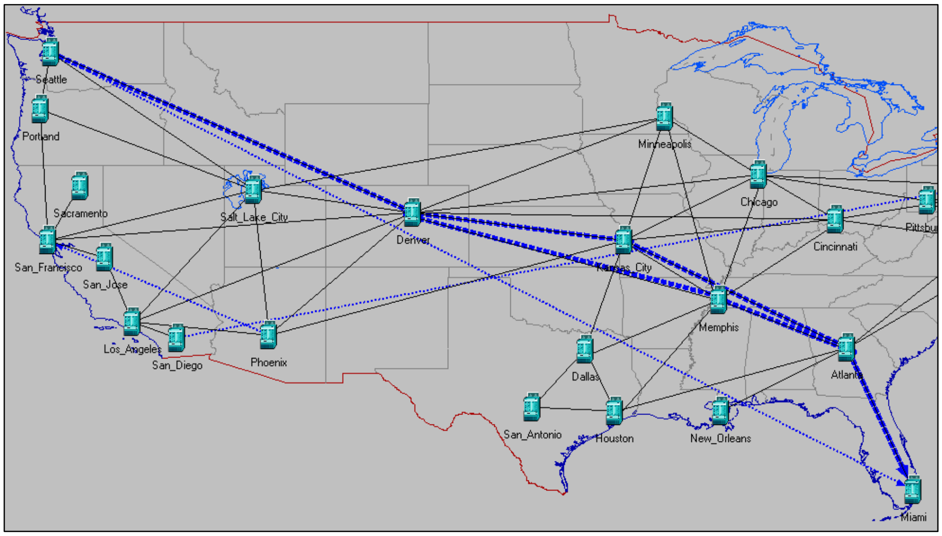

Figure 2-21 Seattle→Miami Route (Between 500 and 1500 Seconds) Displayed

3.7 Click again on the Display field for the selected time to hide the route.

3.8 Click on the corresponding Display field to display the routes taken by the **Seattle→Miami** traffic flow between 1500 and 2000 seconds. Notice that again two equal cost routes are available—one via **Memphis** and one via **Kansas_City**. This happens because the **Kansas_City** router recovered at 1500 seconds.

3.9 Click the **Clear All Routes** button in the Route Report for IP Traffic Flows dialog box.

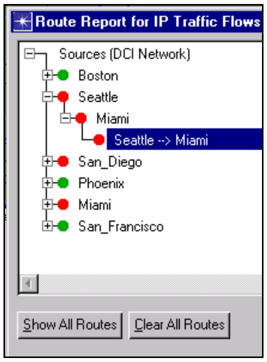

Route Report for IP Traffic Flows

Sources (DCI Network)
- Boston
- Seattle
  - Miami
    - Seattle --> Miami
- San_Diego
- Phoenix
- Miami
- San_Francisco

Figure 2-22 Route Report for IP Traffic Flows dialog box

3.10 Click on the Display field corresponding to routes taken after 2000 seconds. Notice that no route is displayed in the network. Because we have configured the link between **Atlanta** and **Miami** to fail after 2000 seconds, the **Miami** destination has become unreachable, and any route to that destination has been flushed from the OSPF routing tables of all routers across the network

3.11 Expand the **San_Diego** source by clicking on the (+) buttons until the name of **San_Diego→Boston** traffic flow is visible. Click on the traffic flow. On the right side, a table of snapshots of routes used for routing the traffic flow at different times is shown.

3.12 Click on the corresponding Display field to display the routes taken by the **San_Diego→Boston** traffic flow at approximately 168 seconds; seven equal cost routes are taken.

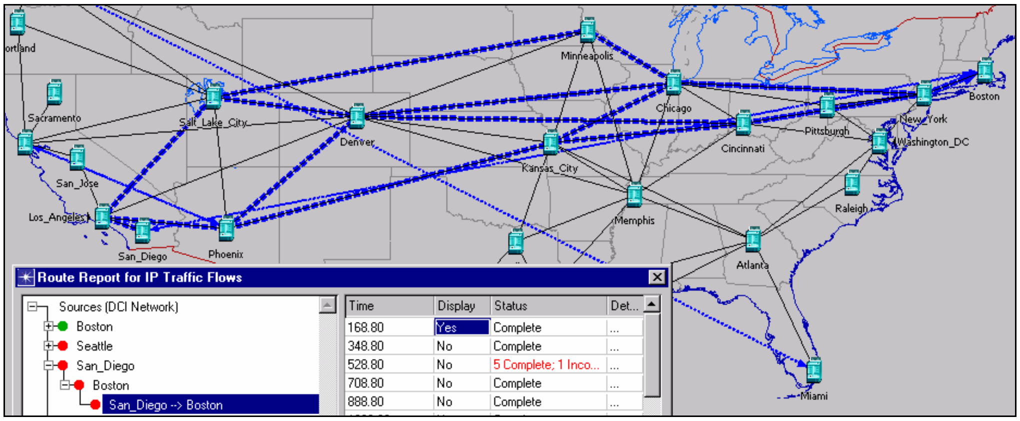

Route Report for IP Traffic Flows

- Sources (DCI Network)
  - Boston
  - Seattle
  - San_Diego
    - Boston
      - San_Diego --> Boston

Time	Display	Status	Det...
168.80	Yes	Complete	...
348.80	No	Complete	...
528.80	No	5 Complete; 1 Inco...	...
708.80	No	Complete	...
888.80	No	Complete	...

Figure 2-23 San_Diego→Boston Route (At 168 Seconds) Displayed

3.13 In the Route Report for IP Traffic Flows dialog box, click **Clear All Routes**.

3.14 Expand the Boston source by clicking on the (+) buttons until the name of the **Boston→San_Diego** traffic flow is visible. Click on the traffic flow name. On the right side, a table of route snapshots taken at different times appears.

3.15 Click on the Display field corresponding to the route taken by the **Boston→San_Diego** traffic flow at approximately 166 seconds; two equal cost paths are taken.

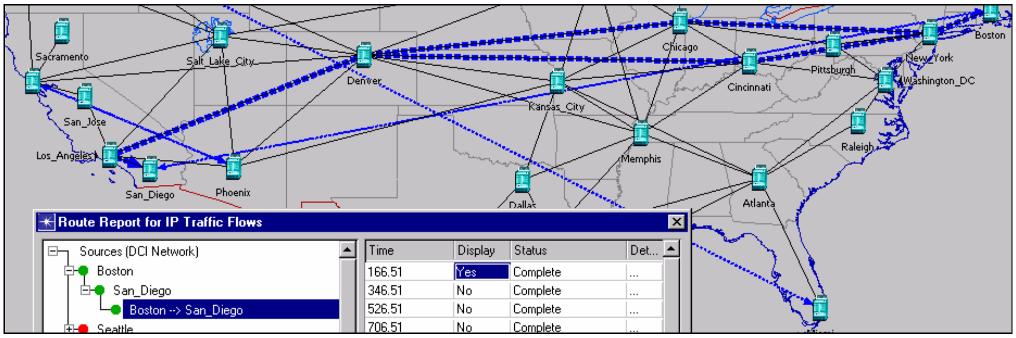

Route Report for IP Traffic Flows

- Sources (DCI Network)
  - Boston
    - San_Diego
      - Boston --> San_Diego
  - Seattle

Time	Display	Status	Det...
166.51	Yes	Complete	...
346.51	No	Complete	...
526.51	No	Complete	...
706.51	No	Complete	...

Figure 2-24 Boston→San_Diego Route (At 166 Seconds) Displayed

3.16 Notice that the traffic from **San Diego** to **Boston** is routed on seven equal cost paths, whereas the reverse traffic from **Boston** to **San Diego** is routed only on two equal cost paths. This is a case of *asymmetric routing*: the traffic is routed in one direction on a set of paths and routed in the reverse direction on a different set of paths. This is not desirable behavior in many cases, especially for applications like Voice over IP (VoIP), since the end-to-end delays in either direction is different.

3.17 Finally, expand the Phoenix source by clicking on the (+) buttons until the name of **Phoenix→San Francisco** traffic flow is visible. Click on the traffic flow name.

3.18 Click in the Display field corresponding to approximately 168 seconds to display the route taken by the **Phoenix→San Francisco** traffic flow.

3.19 Note that the demand goes from **Phoenix** → **Denver** → **San Francisco**.

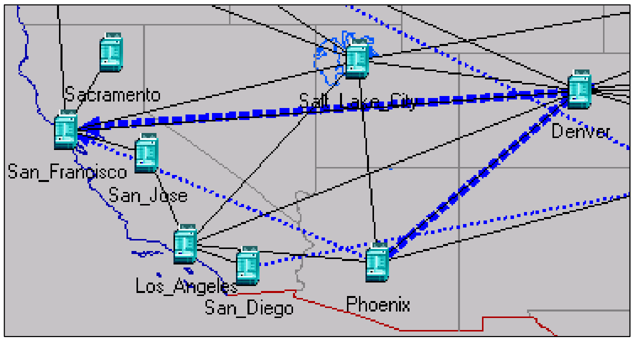

Figure 2-25 Phoenix→San Francisco Route (At 168 Seconds) Displayed

3.20 Finally, in the Route Report for IP Traffic Flows dialog box, click **Clear All Routes**. Click **Close** to dismiss this window.

4 View routing convergence statistics.

4.1 Right-click on the project space (away from all objects), and choose **View Results**.

4.2 Expand the OSPF category under Global Statistics by clicking on the (+) button.

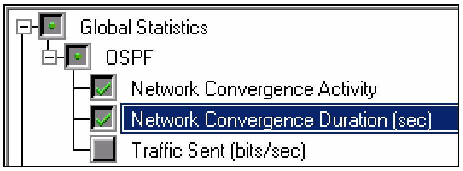

Global Statistics
OSPF
☑ Network Convergence Activity
☑ Network Convergence Duration (sec)
□ Traffic Sent (bits/sec)

Figure 2-26 Statistics Selected in View Results dialog box

4.3 Check the **Network Convergence Activity** and **Network Convergence Duration** statistics; then click **Show**.

4.4 To change the time scale to seconds, right-click on the edge of the graph panel, and choose **Time Axis > Seconds**.

4.5 Zoom in to the relevant portion of the graph by highlighting the four graph points.

To view the coordinates of a point in the statistics graph, place the mouse tip in the neighborhood of the point. A box like the one below will display highlighting the coordinates of the point on the graph nearest to the mouse tip.

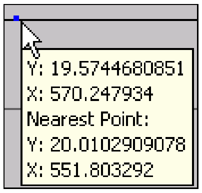

Y: 19.5744680851
X: 570.247934
Nearest Point:
Y: 20.0102909078
X: 551.803292

Figure 2-27 Graph Coordinates Annotation

4.6 The convergence statistics panel displays, as shown below.

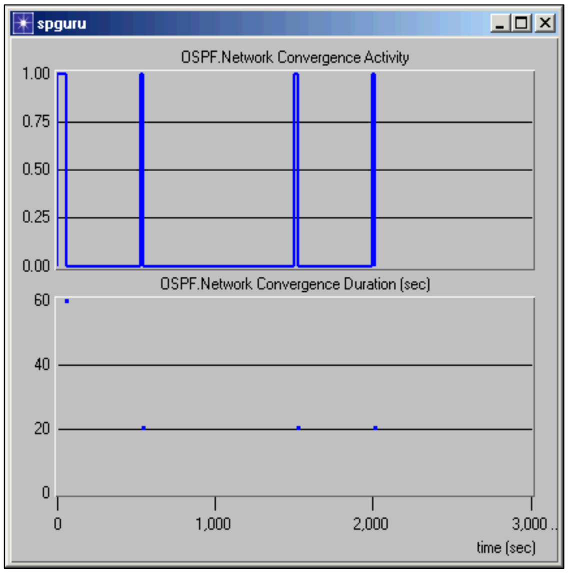

Figure 2-28 Sample Convergence Statistics Graph

The **OSPF.Network Convergence Activity** graph plots the value **1** for the time interval during which convergence activity is detected in the network, i.e. when LSAs are being sent and received. It plots the value **0** at other times. The **OSPF.Network Convergence Duration** plots a data point that indicates the time it took the network to converge after each period of convergence activity. The value on the x-axis corresponds to the moment in time convergence was achieved, while the y-axis value corresponds to the actual convergence time.

From the **OSPF.Network Convergence Duration** graph, notice that the convergence duration is longest when the network ramps-up—it lasts 60 seconds. After each of the subsequent failure or recovery events, convergence is achieved in about 20 seconds.

4.7 For the convergence activity due to the failure of node **Kansas_City** at 500 seconds, point the mouse tip to the top of the corresponding “activity peak” in the **OSPF.Network Convergence Activity** graph (if needed, click to bring the graph window in focus). An annotation will display stating the x-coordinate of the point at the start of the convergence interval. Its x-coordinate corresponds to the time convergence was detected in the network, i.e. 530 seconds, approximately.

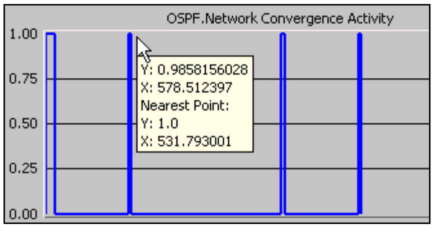

Figure 2-29 Start of Network Convergence Activity

4.8 Similarly, point the mouse tip at the base of the “activity peak” around 500 seconds; the x-coordinate of the nearest point in the annotation gives the time when the network converged, which is approximately 550 seconds. It is interesting to notice that although failure occurred at 500 seconds, it was not detected by the protocol until 30 seconds afterwards, leading to convergence at about 50 seconds after the router went down.

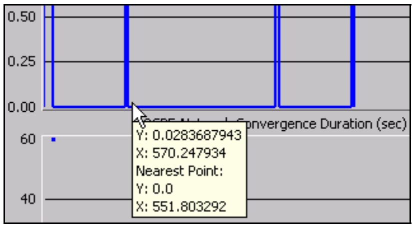

Figure 2-30 Network Converged

4.9 Repeat the last two steps for the “activity peak” following the recovery of **Kansas_City** at 1500 seconds; notice that convergence activity started at approximately 1510 seconds and finished at approximately1530 seconds.

4.10 Close the graph panel by clicking the **X** at the top-right of the graph window.

5 View routing table size statistics.

5.1 Back in the View Results dialog box, click **Unselect** to unselect previously selected graphs.

5.2 Expand the DCI Network category under Object Statistics by clicking on the (+) button.

5.3 Expand **San Diego** by clicking on the (+) button. Similarly expand the **Route Table Category**; check **Size (number of entries)**, and click **Show**.

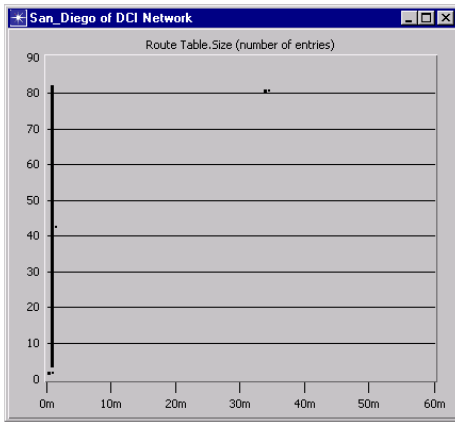

Figure 2-31 San Diego DCI Network Route Table Size

5.4 The graph displays the average number of entries, next to a bar extending from the minimum (on the y-axis) to the maximum number of entries (on the y-axis) over a bucket period; notice that the routing table for San Diego has as many as 82 entries, with more variability when the table is built initially and less variability afterwards. Notice the vertical line between the minimum and the maximum shrinks to almost one point after 30 minutes, when only a minor change occurs in the table.

End of Procedure 2-4

Conclusion

We examined some routing results on a flat OSPF network. We saw that in such a network the IP Forwarding Tables are relatively large with over 80 entries. We also noticed that the Link State Database on all the routers is the same. All routers have 26 Router Links LSAs and no LSAs of any other type. We also saw that given a traffic flow, OSPF can simultaneously route it through different equal cost paths, providing a degree of flexibility in case of failure. Finally, we saw that the network takes longest to converge when it ramps up at time 0.

Modifications to Base Network

In this section, we introduce some modifications to the base OSPF network from the previous section, with the general aim of improving performance.

We examine the cause of asymmetric routing, and perform basic traffic engineering to modify the routes. We also examine the influence of SPF timers on network convergence time.

Procedure 2-5 Control Routes Used by Traffic Flows

1 Duplicate the scenario.

1.1 Choose **Scenario > Duplicate Scenario…**.

1.2 Name the new scenario `Lab1b`.

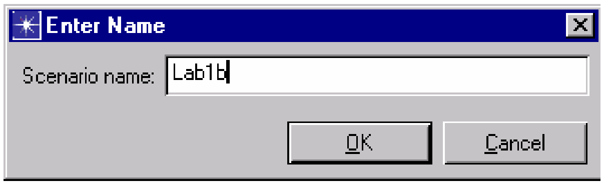

Enter Name

Scenario name: Lab1b

Figure 2-32 Enter Name dialog box

2 Click **OK**.

Remember that the traffic flow between **San Diego** and **Boston** is initially routed via seven different equal-cost paths. We would like to force the traffic to take the link between **Los Angeles** and **Denver**, instead of the two links **Los Angeles↔Salt Lake City** and **Los Angeles↔Phoenix**. A visualization of the seven paths taken by the traffic flow before any failure occured in `Lab1a` is depicted below.

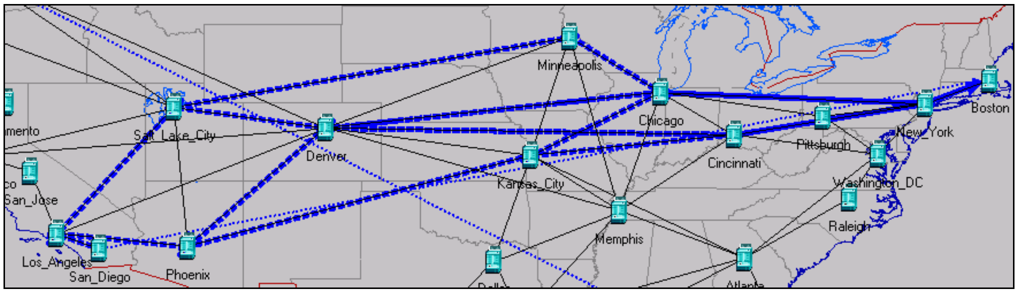

Figure 2-33 Network Model Showing 7 Paths From San Diego to Boston

At the same time, recall that the traffic from **Boston** to **San Diego** was taking only two paths, as shown below.

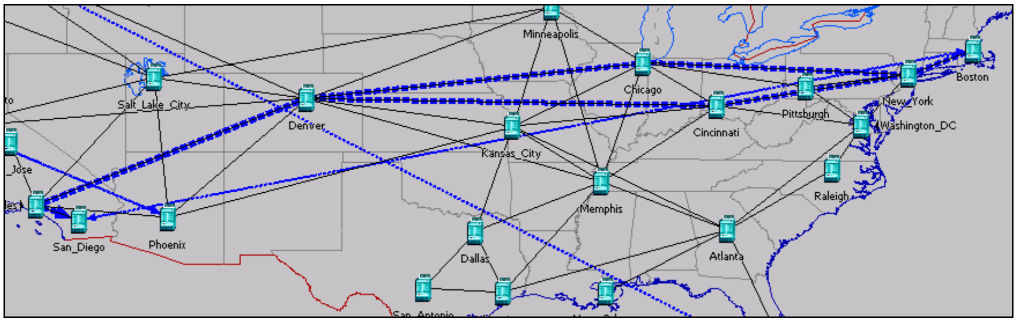

Figure 2-34 Network Model Showing 2 Paths From Boston to San Diego

We referred to this situation in which traffic between the same two nodes takes different paths depending on the direction as *asymmetric routing.* Since all the links in the topology have the same rate, the cause for asymmetric routing appears puzzling. We suspect the cause to be a misconfiguration. To identify a possible misconfiguration, we will run NetDoctor.

3 Run NetDoctor to find the cause of asymmetric routing.

3.1 Choose **NetDoctor > Configure/Run NetDoctor**.

3.2 Select **OSPF** and **IP Routing** from the Rule Suites list on the Rules tab.

3.3 Click **Run**.

3.4 NetDoctor automatically launches a report about the network with the following summary; notice there are no errors, but there are two warnings.

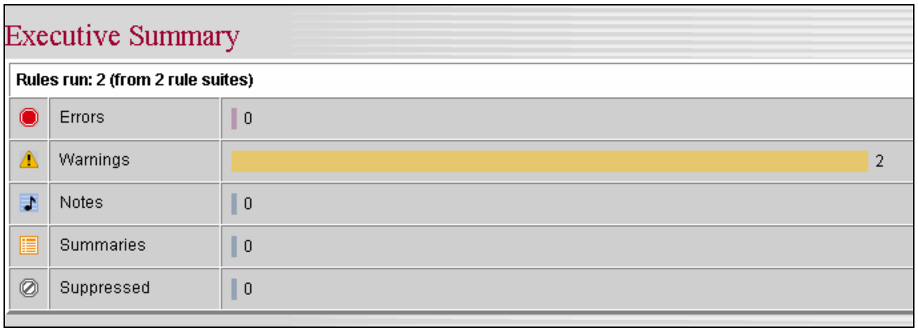

Executive Summary

Rules run: 2 (from 2 rule suites)	
Errors	0
Warnings	2
Notes	0
Summaries	0
Suppressed	0

Figure 2-35 NetDoctor Report

3.5 The warnings are displayed on the left panel; click on the warning titled **OSPF Inconsistent Metric**. A detailed description of the circumstances that led to this warning appears in the right panel.

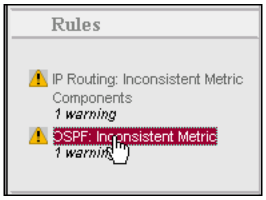

Rules

IP Routing: Inconsistent Metric Components
1 warning

OSPF: Inconsistent Metric
1 warning

Figure 2-36 OSPF Inconsistent Metric Warning

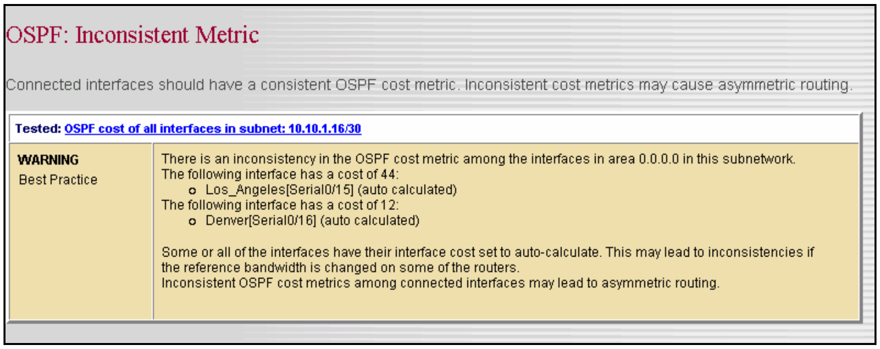

OSPF: Inconsistent Metric

Connected interfaces should have a consistent OSPF cost metric. Inconsistent cost metrics may cause asymmetric routing.

Tested: OSPF cost of all interfaces in subnet: 10.10.1.16/30

WARNING
Best Practice

There is an inconsistency in the OSPF cost metric among the interfaces in area 0.0.0.0 in this subnetwork.
The following interface has a cost of 44:
- Los_Angeles[Serial0/15] (auto calculated)

The following interface has a cost of 12:
- Denver[Serial0/16] (auto calculated)

Some or all of the interfaces have their interface cost set to auto-calculate. This may lead to inconsistencies if the reference bandwidth is changed on some of the routers.
Inconsistent OSPF cost metrics among connected interfaces may lead to asymmetric routing.

Figure 2-37 Details of OSPF Inconsistent Metric Warning

3.6 Notice that the inconsistency reported in the warning can result in asymmetric routing. This makes it relevant to our investigation of asymmetric routing. The warning highlights the fact that OSPF costs on the two interfaces at either side of the **Los Angeles <->Denver** link differ from each other. In addition, it is mentioned that the cost is auto-calculated, i.e. it depends on the bandwidth value of the IP interface. Thus, we will need to look into the configuration of the relevant IP interface in order to understand the metric value. Another warning exists that may shed light on the IP interface configuration.

3.7 Click on the warning titled **IP Routing Inconsistent Metric Components**; a detailed description of the circumstances that led to this warning appears in the right panel.

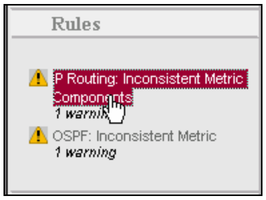

Rules

⚠ IP Routing: Inconsistent Metric Components
1 warning

⚠ OSPF: Inconsistent Metric
1 warning

Figure 2-38 IP Routing Inconsistent Metric Components Warning

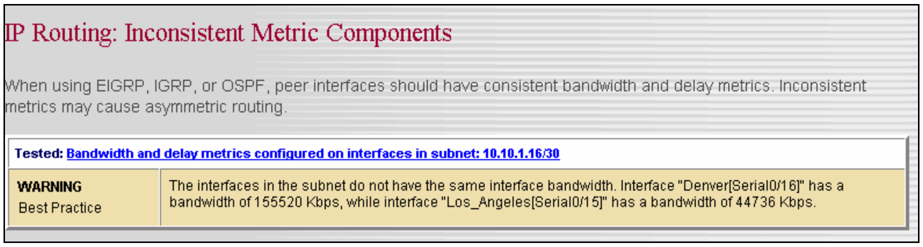

IP Routing: Inconsistent Metric Components

When using EIGRP, IGRP, or OSPF, peer interfaces should have consistent bandwidth and delay metrics. Inconsistent metrics may cause asymmetric routing.

Tested: Bandwidth and delay metrics configured on interfaces in subnet: 10.10.1.16/30	
WARNING Best Practice	The interfaces in the subnet do not have the same interface bandwidth. Interface "Denver[Serial0/16]" has a bandwidth of 155520 Kbps, while interface "Los_Angeles[Serial0/15]" has a bandwidth of 44736 Kbps.

Figure 2-39 Detail of IP Routing Inconsistent Metric Components Warning

3.8 The warning indicates that the interface bandwidth (used to calculate the OSPF metric) is different on the two interfaces. Since the link between **Los Angeles** and **Denver** is symmetric (i.e. has the same data rate on both directions), there is no immediate reason why the interface bandwidth should be different on the two interfaces. To correct this configuration, we can modify one interface bandwidth and make it equal to the other interface. In order to make the **Los Angeles <->Denver** link preferable over the other links, we will pick the higher of the two bandwidths, i.e. 155.5 Mbps. Hence we will change **Los Angeles'** interface Serial0/15 to a bandwidth of 155.5 Mbps.

3.9 Right-click on the **Los Angeles** router and choose **Edit Attributes**.

3.10 Click on the (+) button adjacent to IP, and choose Interface Information under IP Routing Parameters; scroll to interface Serial0/15, and double-click in the Metric Information corresponding to this interface.

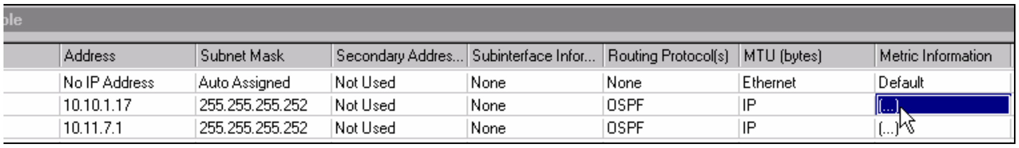

ble

Address	Subnet Mask	Secondary Addres...	Subinterface Infor...	Routing Protocol(s)	MTU (bytes)	Metric Information
No IP Address	Auto Assigned	Not Used	None	None	Ethernet	Default
10.10.1.17	255.255.255.252	Not Used	None	OSPF	IP	(...)
10.11.7.1	255.255.255.252	Not Used	None	OSPF	IP	(...)

Figure 2-40 Interface Information Table

3.11 Double-click in the Value column corresponding to the attribute to edit it. Change the existing value of 44,736 to 155,520; keep clicking **OK** to dismiss all open windows.

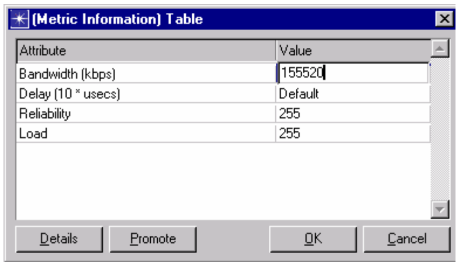

(Metric Information) Table

Attribute	Value
Bandwidth (kbps)	155520
Delay (10 * usecs)	Default
Reliability	255
Load	255

Figure 2-41 Metric Information Table

4 Run NetDoctor:

4.1 Choose **NetDoctor > Configure/Run NetDoctor**.

4.2 Click **Run**.

4.3 Notice that NetDoctor no longer reports the asymmetric routing warnings.

5 Change OSPF timers to improve convergence.

5.1 Choose **Protocols > OSPF > Configure SPF Calculations Paramters...**.

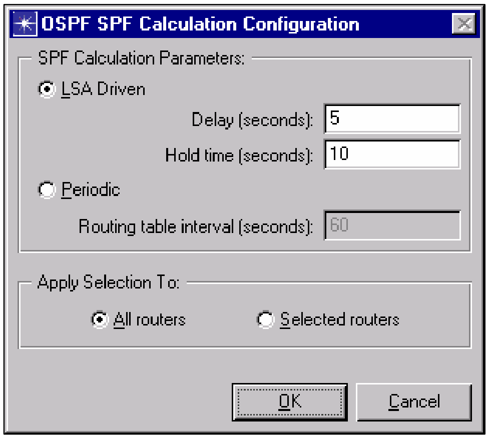

OSPF SPF Calculation Configuration

SPF Calculation Parameters:

● LSA Driven

Delay (seconds): 5

Hold time (seconds): 10

○ Periodic

Routing table interval (seconds): 60

Apply Selection To:

● All routers ○ Selected routers

Figure 2-42 OSPF SPF Calculation Configuration dialog box

5.2 Make sure the **All routers** radio button is selected.

5.3 Enter **5** for Delay and **10** for Hold time.

5.4 Click **OK**.

6 Run the simulation.

6.1 Click the **Configure/Run Simulation** toolbar button.

6.2 Click **Run** to start the simulation.

6.3 Close the Simulation Sequence dialog box after the simulation runs.

End of Procedure 2-5

Procedure 2-6 Visualize Route Changes

1 View traffic-engineering results.

1.1 Choose **Protocols > IP > Demands > Display Routes for Configured Demands...**.

1.2 Expand the **San Diego** source by clicking on the (+) buttons, until the name of the traffic flow **San Diego** → **Boston** appears. Click on the traffic flow name.

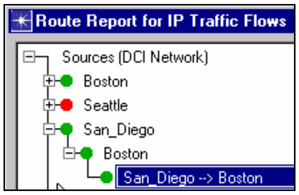

Route Report for IP Traffic Flows

- Sources (DCI Network)
  - ● Boston
  - ● Seattle
  - ● San_Diego
    - ● Boston
      - ● San_Diego --> Boston

Figure 2-43 Route Report for IP Traffic Flows dialog box

1.3 Click on the corresponding Display field to display the route taken by the traffic flow at approximately 168 seconds; notice that the traffic flow is routed through the link **Los Angeles** ↔ **Denver** – in other words, by changing the bandwidth on an IP interface, we achieved a traffic-engineering objective.

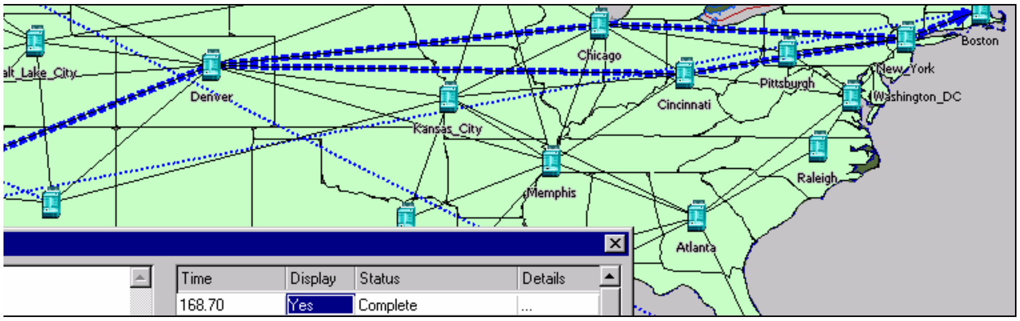

Time	Display	Status	Details
168.70	Yes	Complete	...

Figure 2-44 Displayed Traffic Flow

2 Compare convergence results.

2.1 Right-click anywhere in the project space and choose **Compare Results**.

2.2 In the left panel, expand the OSPF category under Global Statistics by clicking the (+) button.

2.3 Check Network Convergence Duration (sec).

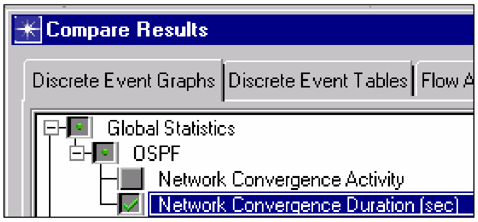

Compare Results

Discrete Event Graphs | Discrete Event Tables | Flow A

- Global Statistics
  - OSPF
    - □ Network Convergence Activity
    - ☑ Network Convergence Duration (sec)

Figure 2-45 Network Convergence Duration Attribute

2.4 In the pull down options at the bottom of the right side panel, select **Overlaid Statistics** and **Select Scenarios**.

2.5 Check **Lab1a** and **Lab1b**, and then click **OK**.

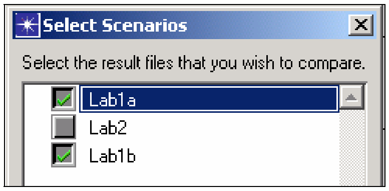

Select Scenarios

Select the result files that you wish to compare.

☑ Lab1a
□ Lab2
☑ Lab1b

Figure 2-46 Select Scenarios dialog box

2.6 Click **Show**.

2.7 Notice that lowering the SPF parameters (Delay and Hold Time) had the effect of lowering the convergence time following a disturbance (e.g. network start-up, failure, or recovery).

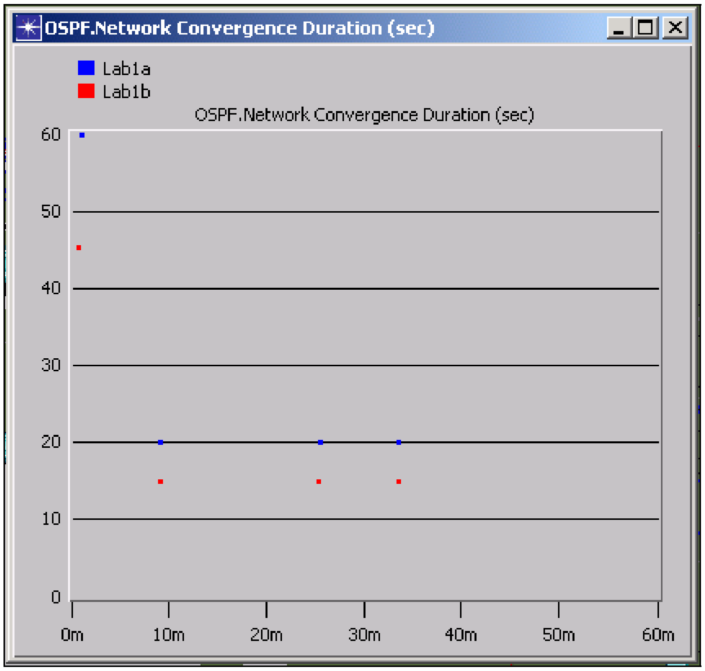

Figure 2-47 OSPF Network Convergence Statistics

2.8 Click **Close** to dismiss the Compare Results window.

End of Procedure 2-6

Conclusion

In the previous section, we saw how traffic engineering of flow routes can be achieved by modifying the cost of an interface. We also saw the effect of lowering the SPF parameters on the convergence duration: the smaller the Delay and Hold Time, the faster the network converges after a disturbance.

OSPF Network with Areas

In this section, an OSPF network with areas has replaced the flat OSPF network of the previous example. We study the effect of defining OSPF areas on the size of IP forwarding tables and OSPF link state databases, as well as on the routes taken by demands.

Procedure 2-7 Configure OSPF with Areas Scenario

1 Switch to new scenario.

1.1 Choose **Scenarios > Switch to Scenario > Lab2**.

2 Visualize areas.

2.1 Choose **View > Visualize Protocol Configuration > OSPF Area Configuration**…

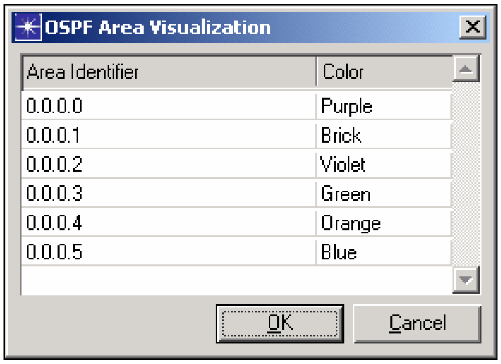

OSPF Area Visualization

Area Identifier	Color
0.0.0.0	Purple
0.0.0.1	Brick
0.0.0.2	Violet
0.0.0.3	Green
0.0.0.4	Orange
0.0.0.5	Blue

Figure 2-48 OSPF Area Visualization dialog box

2.2 The dialog box prompts you to accept colors for each OSPF area. Click **OK**.

The links in each area will appear colored with the color corresponding to the area.

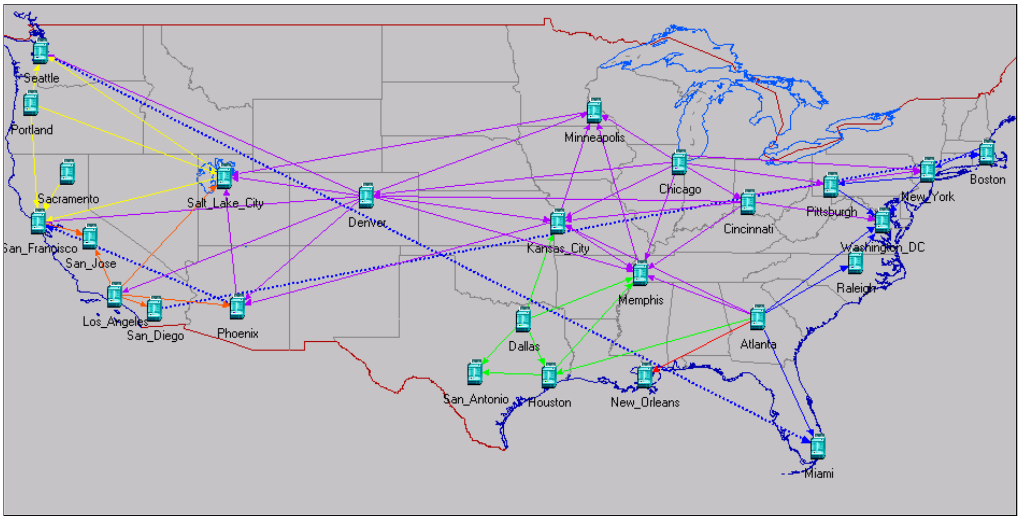

Figure 2-49 Color-coded Links

2.3 Notice the link between **Atlanta** and **Miami** is colored blue, meaning that it belongs to Area 5. We would like to change this link to Area 4 (i.e., its connected OSPF interfaces will belong to Area 4); to do that, select the link by clicking on it once, then choose **Protocols > OSPF > Configure Areas**.

2.4 In the OSPF Area Configuration, enter **0.0.0.4** for Area Identifier; click **OK**.

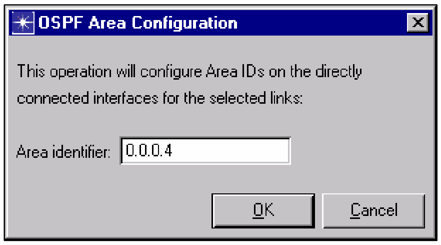

OSPF Area Configuration

This operation will configure Area IDs on the directly connected interfaces for the selected links:

Area identifier: 0.0.0.4

Figure 2-50 OSPF Area Configuration dialog box

2.5 Press **Ctrl+Shift+C** to clear the OSPF area visualization.

2.6 Again choose **View > Visualize Protocol Configuration > OSPF Area Configuration**; click **OK** to confirm the colors.

You should see the link between Atlanta and Miami colored orange, confirming that it is now within Area 4, as configured.

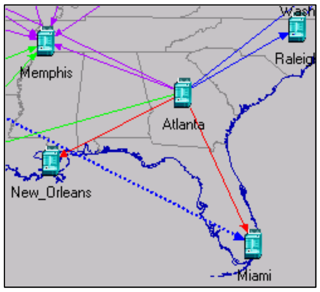

Figure 2-51 Link Between Atlanta and Miami

3 Configure OSPF summaries.

3.1 Notice that the **Atlanta** router has interfaces in 4 areas: 0, 3, 4 and 5.

3.2 We would like to advertise summary addresses for all 4 areas meeting on router **Atlanta**. To that end, examine the Area Summarization Table of router **Atlanta**. Right-click on the **Atlanta** router and choose **Edit Attributes**.

3.3 Expand the **IP Routing Protocols** category by clicking on the adjacent (+) button; click in the **Value** column next to **OSPF Parameters**.

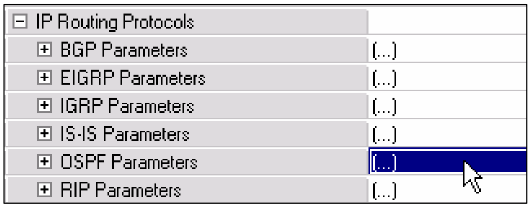

⊟ IP Routing Protocols
- ⊞ BGP Parameters (...)
- ⊞ EIGRP Parameters (...)
- ⊞ IGRP Parameters (...)
- ⊞ IS-IS Parameters (...)
- ⊞ OSPF Parameters (...)
- ⊞ RIP Parameters (...)

Figure 2-52 IP Routing Protocols Attributes dialog box

3.4 Double-click in the **Value** field next to **Processes.**

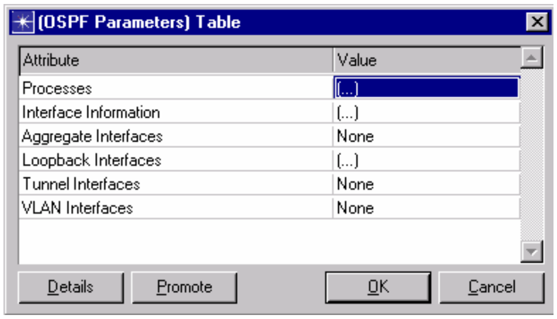

(OSPF Parameters) Table

Attribute	Value
Processes	(...)
Interface Information	(...)
Aggregate Interfaces	None
Loopback Interfaces	(...)
Tunnel Interfaces	None
VLAN Interfaces	None

Figure 2-53 OSPF Parameters Table dialog box

3.5 Double-click in the **Process Parameters** field.

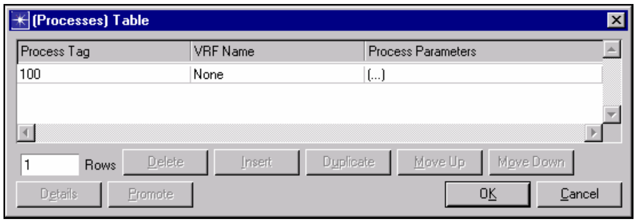

Process Tag	VRF Name	Process Parameters
100	None	(...)

Rows: 1

Figure 2-54 Processes Table dialog box

3.6 Double-click in the **Value** field next to **Area Summarization.**

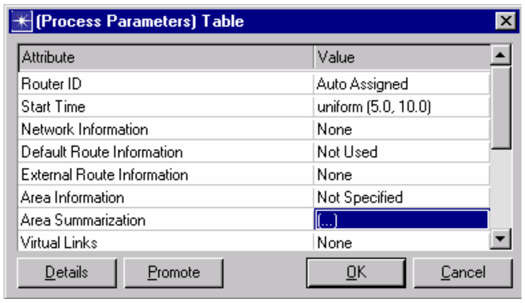

Attribute	Value
Router ID	Auto Assigned
Start Time	uniform (5.0, 10.0)
Network Information	None
Default Route Information	Not Used
External Route Information	None
Area Information	Not Specified
Area Summarization	(...)
Virtual Links	None

Figure 2-55 Process Parameters Table dialog box

3.7 In the table of area summarizations, notice that only two areas are summarized: area 0 and area 5. We will introduce similar summaries for areas 4 and 3.

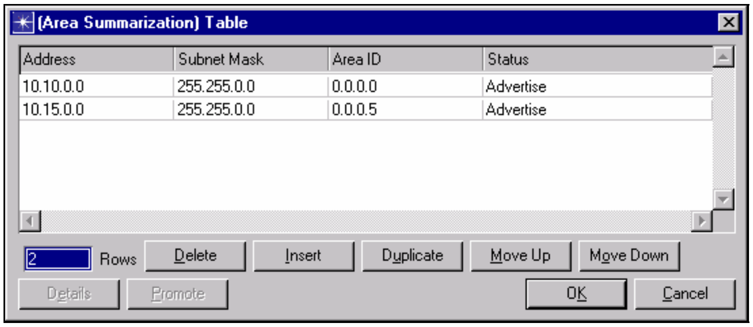

Address	Subnet Mask	Area ID	Status
10.10.0.0	255.255.0.0	0.0.0.0	Advertise
10.15.0.0	255.255.0.0	0.0.0.5	Advertise

Rows: 2

Figure 2-56 Area Summarization Table dialog box

3.8 To add two more rows to the table, click on the box specifying the number of rows in the lower left corner of the window and choose **Edit…** Enter **4** for **Rows**. Click on the window panel outside the **Rows** box: two more rows will appear in the table.

3.9 Fill in the two new rows with information as shown below.

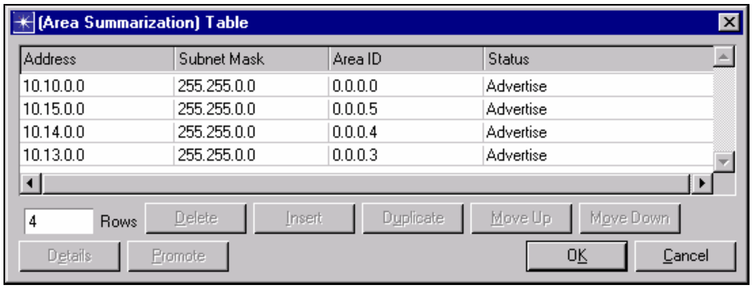

Address	Subnet Mask	Area ID	Status
10.10.0.0	255.255.0.0	0.0.0.0	Advertise
10.15.0.0	255.255.0.0	0.0.0.5	Advertise
10.14.0.0	255.255.0.0	0.0.0.4	Advertise
10.13.0.0	255.255.0.0	0.0.0.3	Advertise

4 Rows

Figure 2-57 Edited Area Summarization Table dialog box

3.10 Click **OK** on each dialog box to save changes.

End of Procedure 2-7

Procedure 2-8 Run Simulation and View Results

1 Run the simulation.

1.1 Click the **Configure/Run Simulation** toolbar button.

1.2 Click **Run** to start the simulation.

1.3 Close the Simulation Sequence dialog box after the simulation runs.

2 Inspect the OSPF LS Database

2.1 Right-click on the **San_Diego** node, and select **View Results**.

2.2 Choose the **Discrete Event Tables** tab.

2.3 Expand the **Performance** category by clicking on the adjacent (+) button.

2.4 Click on **OSPF Link State Database Summary at 3600 seconds**.

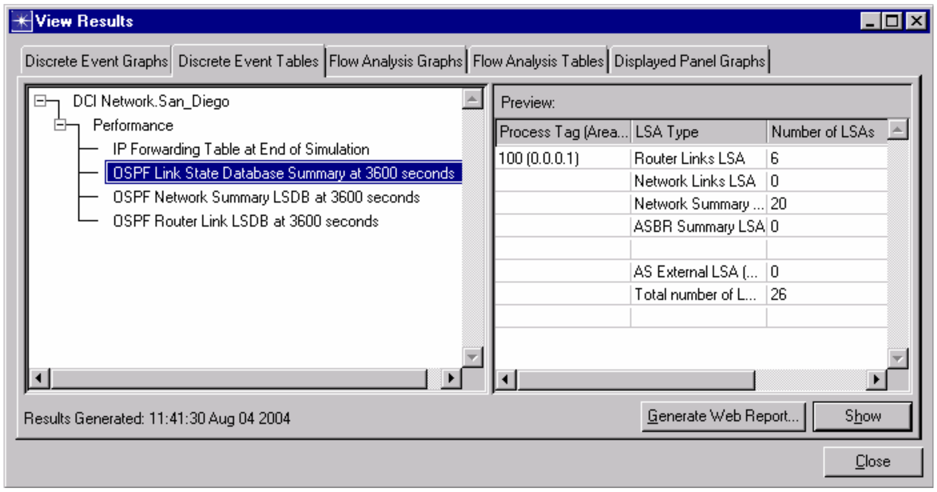

DCI Network.San_Diego
- Performance
  - IP Forwarding Table at End of Simulation
  - OSPF Link State Database Summary at 3600 seconds
  - OSPF Network Summary LSDB at 3600 seconds
  - OSPF Router Link LSDB at 3600 seconds

Process Tag (Area...	LSA Type	Number of LSAs
100 (0.0.0.1)	Router Links LSA	6
	Network Links LSA	0
	Network Summary ...	20
	ASBR Summary LSA	0
	AS External LSA (...	0
	Total number of L...	26

Results Generated: 11:41:30 Aug 04 2004

Figure 2-58 Discrete Event Tables Tab for San_Diego Node

2.5 Notice that **San_Diego**, which is an area-internal router (i.e. belongs to only one area), has an LSDB for area 0.0.0.1. It has only 6 Router Links LSAs in its LSDB, and 20 Network Summary LSAs.

2.6 Click **Close** to dismiss the View Results window.

2.7 Similarly, right-click on the **Atlanta node** and select **View Results**.

2.8 Choose the **Discrete Event Tables** tab. Expand the **Performance** category; click on the **OSPF Link State Database Summary at 3600 seconds**.

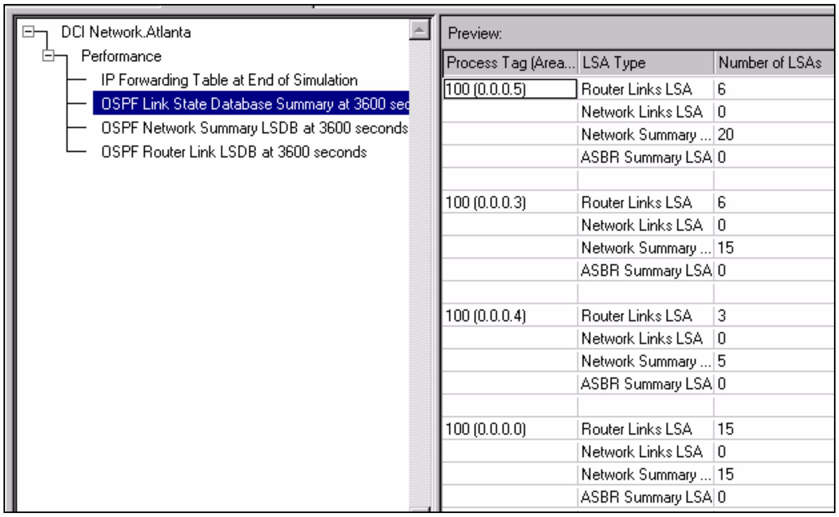
DCI Network.Atlanta
Performance
IP Forwarding Table at End of Simulation
OSPF Link State Database Summary at 3600 sec
OSPF Network Summary LSDB at 3600 seconds
OSPF Router Link LSDB at 3600 seconds
Preview:

Process Tag (Area...	LSA Type	Number of LSAs
100 (0.0.0.5)	Router Links LSA	6
	Network Links LSA	0
	Network Summary ...	20
	ASBR Summary LSA	0
100 (0.0.0.3)	Router Links LSA	6
	Network Links LSA	0
	Network Summary ...	15
	ASBR Summary LSA	0
100 (0.0.0.4)	Router Links LSA	3
	Network Links LSA	0
	Network Summary ...	5
	ASBR Summary LSA	0
100 (0.0.0.0)	Router Links LSA	15
	Network Links LSA	0
	Network Summary ...	15
	ASBR Summary LSA	0

Figure 2-59 Discrete Event Tables Tab for Atlanta Node

2.9 Notice that **Atlanta** has a very different LSDB than **San_Diego**.

2.10 **Atlanta** has separate LSDBs for the four areas (0.0.0.0, 0.0.0.3, 0.0.0.4, and 0.0.0.5) that it is connected to. The number of router links LSAs and Network Summary LSAs in each area is different.

2.11 Click **Close** to dismiss the View Results window.

3 View routing table size statistic.

3.1 Right-click somewhere in the project space (away from all objects).

3.2 Choose **View Results**.

3.3 Expand the **DCI Network** category by clicking on the adjacent (+) button.

3.4 Similarly, expand **San_Diego**; expand the **Route Table**. Click on **Size (number of entries)** statistic to select it; click **Show**.

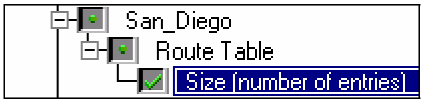

Figure 2-60 Show San_Diego Route Table

3.5 Zoom in the graph traces; notice that the routing table size of **San_Diego** reaches a maximum of 16 entries, a marked improvement in size over the maximum of 82 entries obtained for the same router in Lab1a.

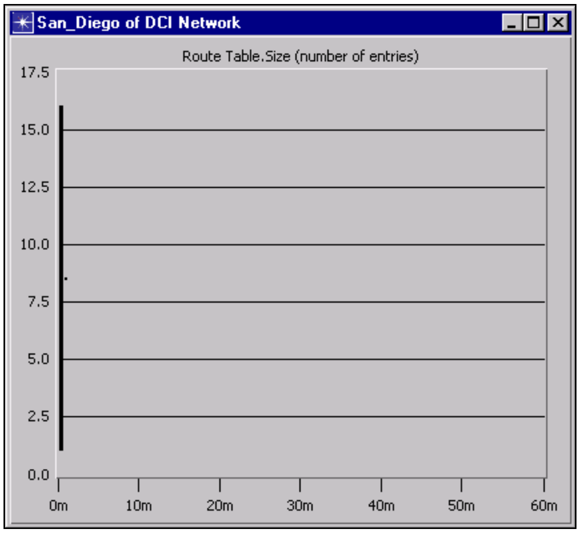

Figure 2-61 San_Diego Route Table Size Graph

3.6 Click on the **Hide/Show Graph Panels** button.

3.7 Click **Close** to dismiss the View Results window.

4 Examine traffic flow routing.

4.1 Choose **Protocols > IP > Demands > Display Routes for Configured Demands....**

4.2 Expand the **Seattle** source by clicking on the (+) buttons twice until the traffic flow **Seattle <-> Miami** appears.

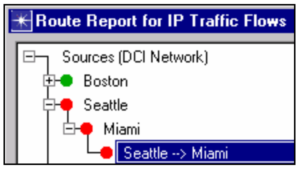

Figure 2-62 Select the Seattle <-> Miami Traffic Flow

4.3 Click on the traffic flow name to view snapshots of the route taken at different times by the traffic flow; these will appear in the right side of the panel. Notice that all the snapshots taken after 2000 seconds will be tagged Incomplete; this is due to failure of the link **Atlanta<->Miami**.

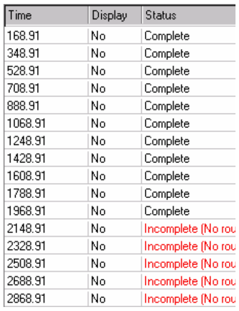

Time	Display	Status
168.91	No	Complete
348.91	No	Complete
528.91	No	Complete
708.91	No	Complete
888.91	No	Complete
1068.91	No	Complete
1248.91	No	Complete
1428.91	No	Complete
1608.91	No	Complete
1788.91	No	Complete
1968.91	No	Complete
2148.91	No	Incomplete (No rou
2328.91	No	Incomplete (No rou
2508.91	No	Incomplete (No rou
2688.91	No	Incomplete (No rou
2868.91	No	Incomplete (No rou

Figure 2-63 Seattle <-> Miami Traffic Flow

4.4 Click in the **Display** column, in the row corresponding to a time after 2000 seconds, to visualize the incomplete route for the traffic flow **Seattle<->Miami**. Notice that this route stops on the **Atlanta** router, which is an area border router for the destination area. If you visualized the same incomplete route in Lab1, inside a flat OSPF network, there will be no segment of the route known, since the destination **Miami** would have been eliminated from all routing tables. However, in this case, even if the destination **Miami** were eliminated from all databases within area 5, there would still be a summary for this destination outside the areas. Consequently, the traffic flow will be routable based on this summarized address up until the border of area 5.

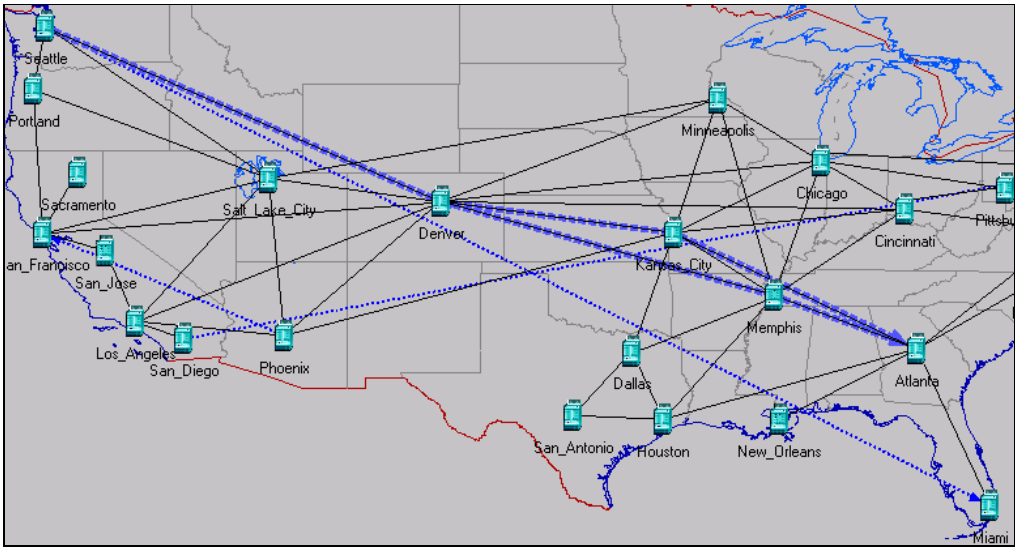

Figure 2-64 Traffic Routes

4.5 Click on **Clear All Routes**, but do not close the Route Report for IP Traffic Flows window.

5 Compare route selection with Lab1b.

5.1 Expand the **San Diego** source by clicking on the (+) buttons until the **San_Diego <-> Boston** traffic flow name appears; click on the traffic flow name to select it.

5.2 Click in the **Display** column, in the row corresponding to time **168.70** seconds to visualize the route taken by the traffic flow **San_Diego <-> Boston**.

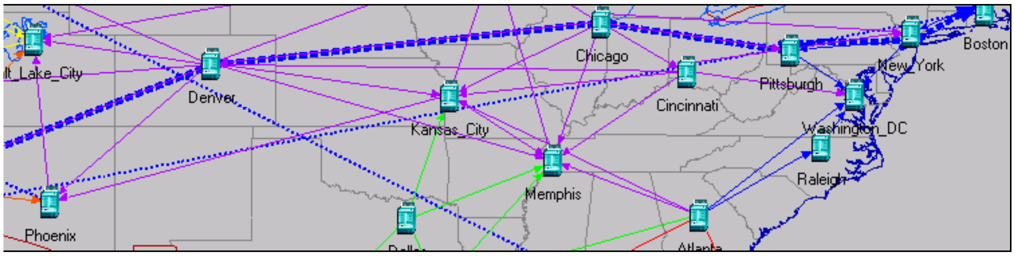

Figure 2-65 Traffic Conditions in Lab2

5.3 Compare with the route taken by the same traffic in the conditions of Lab1.

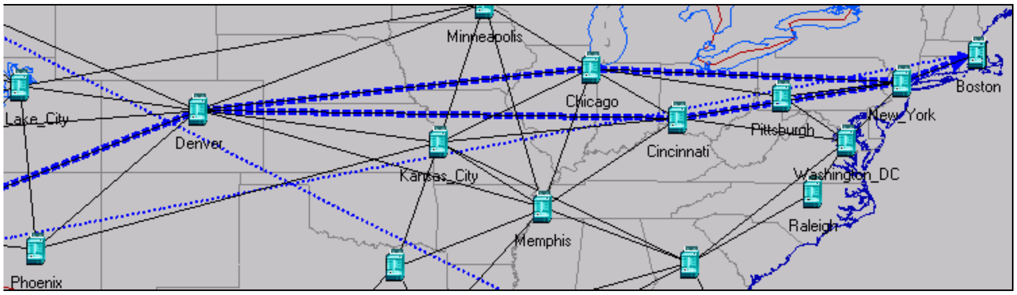

Figure 2-66 Traffic Conditions in Lab1

5.4 Notice that in Lab1, there are 2 equal cost paths between **San Diego** and **Boston**. The path from Lab1 that most closely resembles the path observed in Lab2 has one hop less than the path in Lab2 (i.e. the path in Lab1 goes from **Chicago** directly to **New York,** while the path in Lab2 goes from **Chicago** to **New York** via **Pittsburgh**). This illustrates the point that the reduced overhead and routing table sizes associated with hierarchical routing has a trade-off in that routes selected are not always optimal. In our case, we have seen a slightly longer route being taken in the network with areas compared with the route taken in the flat network between the same source and destination.

5.5 Click on **Clear All Routes**, but do not close the Route Report for IP Traffic Flows window.

5.6 Expand the **Phoenix** traffic source by clicking on the (+) buttons until the **Phoenix <-> San Francisco** traffic flow name shows up; click on the traffic flow name to select it.

5.7 In the right side panel, click on the **Display** field corresponding to approximately **168 seconds**. The following route is displayed in the network for the **Phoenix <-> San Francisco** traffic flow.

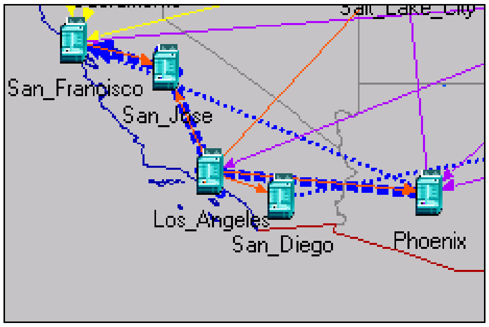

Figure 2-67 Phoenix <-> San Francisco Route

5.8 Compare with the route for the same traffic flow from Lab1a, where areas were not configured. Notice that in the case of areas, a route between the source and destination nodes that stays within one area is preferred to a more optimal route (with regards to number of hops and bandwidth) that crosses different areas. This is another situation where configuring OSPF areas leads to non-optimal routes.

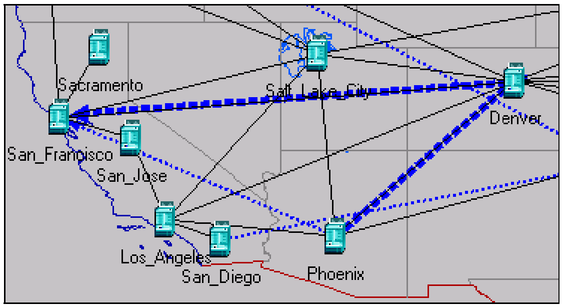

Figure 2-68 Route from Lab1a

End of Procedure 2-8

Conclusion

In this section, we studied the effects of configuring areas in the flat OSPF network of Lab1. We learned that area configuration, together with summarization, leads to smaller routing table sizes and smaller OSPF link state databases. However, we also noticed that area configuration and address summarization also lead to sub-optimal routes chosen by OSPF.

IS-IS Network

In this section, we work with a countrywide IS-IS network, with a backbone and 5 areas. We examine the influence of merging two areas on routes between a source and a destination in these two areas. In doing so, we familiarize you with some of the IS-IS tools available in NETWARS.

Procedure 2-9 Examine the IS-IS Network

1 Switch to the Lab3 scenario by choosing **Scenario > Switch to Scenario > Lab3**.

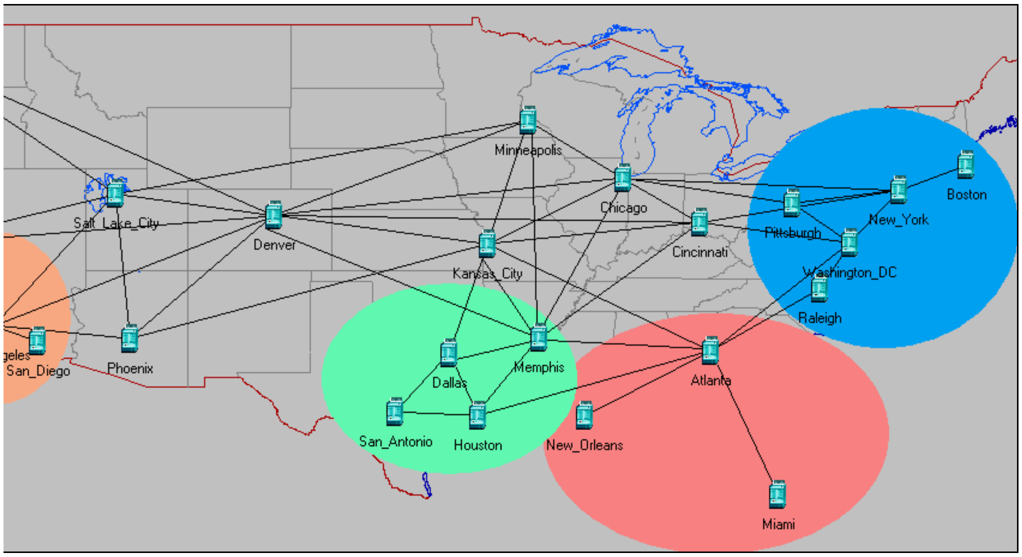

Figure 2-69 Sample Network Scenario

2 Check that the network is running IS-IS on all interfaces.

2.1 Select **View > Visualize Protocol Configuration > IP Routing Domains**.

2.2 Notice that all the links in the network are annotated with the letter “I” signifying that IS-IS is running on the interfaces at either end of each link.

2.3 Press **Ctrl+Shift+C** to clear the routing protocol visualization.

End of Procedure 2-9

Procedure 2-10 Examine the Route Between Nodes in Different Areas

1 Run the simulation.

1.1 Click the **Configure/Run Simulation** toolbar button.

1.2 Click **Run** to start the simulation.

1.3 Close the Simulation Sequence dialog box after the simulation runs.

2 View route between nodes in different areas.

2.1 The **San Jose** router and the **Portland** router are in different areas; we will visualize the routes available in the IP forwarding tables between these nodes.

2.2 Click to select **San Jose**, and then shift-click to also select **Portland**.

2.3 Select **Traffic > Visualize > Open Flows Browser**.

2.4 Click **Show Selected**.

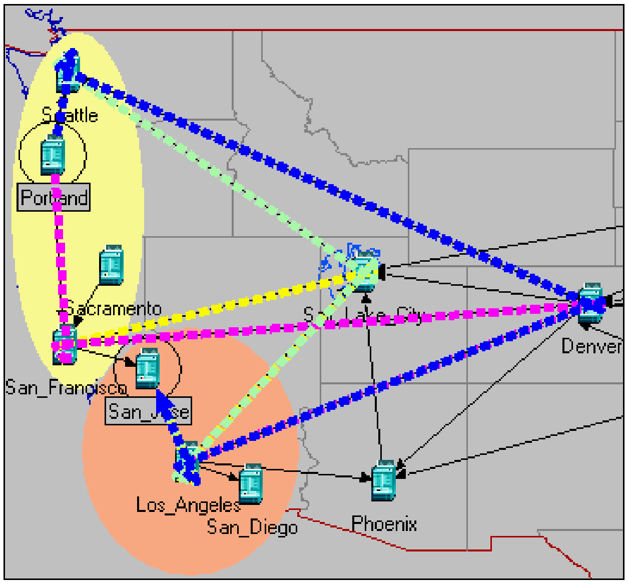

Figure 2-70 View Routes Between Nodes

2.5 We see several routes between **Portland** and **San_Jose**, all of which go through either **Seattle** or **San_Francisco**, which are the border routers for that area. The routes then traverse through the backbone to get to **Los_Angeles** (which is the only border router in **San_Jose's** area), and then finally to **San_Jose**.

2.6 Next we will merge areas **49.0001** (**San_Jose's** area) and **49.0002** (**Portland's** area). Keeping in mind that intra-area routes are preferred over inter-area routes, we will then examine the new route from **Portland** to **San_Jose** and see it follow a path that does not go through the backbone.

End of Procedure 2-10

Procedure 2-11 Merge Two Areas into One

1 Verify current area configuration.

1.1 There are 5 areas configured in the network. To facilitate visualization, colored ovals have been manually added to the network background. The **yellow** oval represents area **49.0002**, while **orange** is area **49.0001**.

1.2 Right-click on the **San Francisco** router, and select **Edit Attributes**.

1.3 Click on the (+) button next to **IP Routing Protocols** to expand it; then click in the **Value** column on the IS-IS Parameters row.

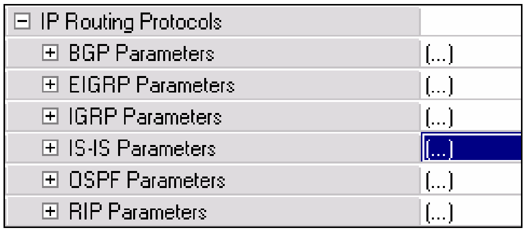

⊟ IP Routing Protocols	
⊞ BGP Parameters	(...)
⊞ EIGRP Parameters	(...)
⊞ IGRP Parameters	(...)
⊞ IS-IS Parameters	(...)
⊞ OSPF Parameters	(...)
⊞ RIP Parameters	(...)

Figure 2-71 IP Routing Protocols Attributes

1.4 Double-click in **Value** field corresponding to **Processes**.

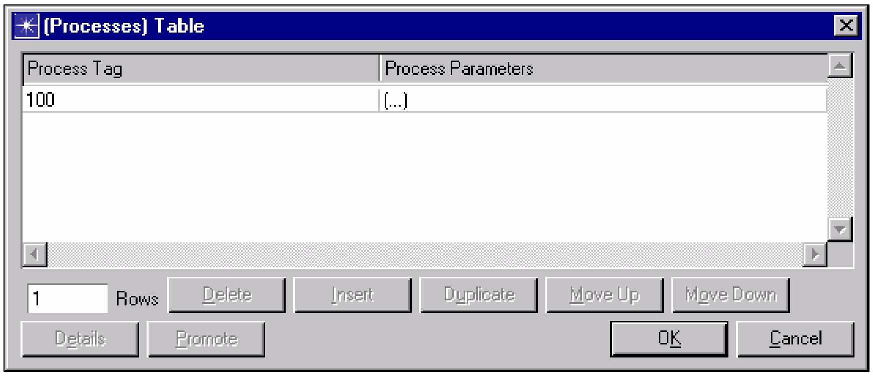

(Processes) Table

Process Tag	Process Parameters
100	(...)

1 Rows

Figure 2-72 Processes Table dialog box

1.5 Double-click in the **Process Parameters** column.

1.6 Double-click in the **Value** field corresponding to **Network Entity Titles**.

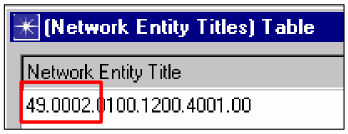

(Network Entity Titles) Table

Network Entity Title
49.0002.0100.1200.4001.00

Figure 2-73 Network Entity Titles Table dialog box

1.7 Notice the prefix **49.0002,** which indicates that **San_Francisco** is in area 49.0002.

2 Merge the two areas.

2.1 Merging areas is very simple in IS-IS. Just add a new NET on the **San Francisco** router with the area ID part of the NET set to **49.0001** (the adjacent area).

In the Network Entity Titles window, click on the first row, which contains the only NET defined on this router, then click **Duplicate**, in the second row. Edit the beginning of the NET to read **49.0001** instead of **49.0002**.

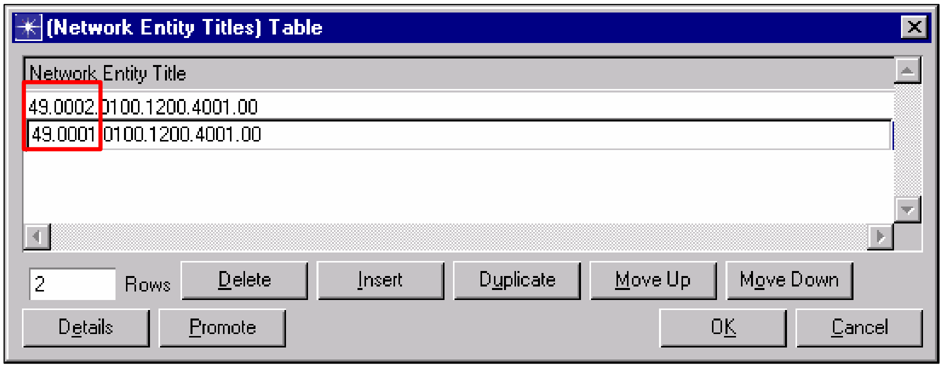

Figure 2-74 Edit the Network Entity Titles Table dialog box

2.2 Click **OK** all the way through to save the changes; the two areas are now merged in one area, simply by configuring both area IDs on **San_Francisco**.

End of Procedure 2-11

Procedure 2-12 Verify that the Areas have been Merged

1 Run the simulation.

1.1 Click the **Configure/Run Simulation** toolbar button.

1.2 Click **Run** to start the simulation.

1.3 Close the Simulation Sequence dialog box after the simulation runs.

2 View the new route between **Portland** and **San_Jose**.

2.1 Click to select **San Jose**, and then shift-click to also select **Portland**.

2.2 Select **Traffic > Visualize > Open Flows Browser**.

2.3 Click **Show Selected**.

2.4 Click **Show All** to display the routes in the network; notice that now there is only one route – **Portland <-> San_Francisco <-> San_Jose**. This route does not go through the backbone, indicating that these two areas have now been merged into one.

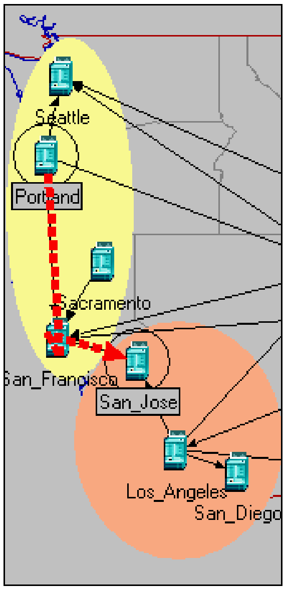

Figure 2-75 Show Merged Routes

End of Procedure 2-12

Conclusion

In this section, we learned that hierarchical routing in IS-IS follows a route selection process similar to that in OSPF, i.e. intra-area routes are preferred over inter-area routes. We also saw how easy it is to merge two IS-IS areas. Trying to merge two areas in OSPF would require a significant amount of reconfiguration.

3 Planning and Analyzing EIGRP and RIP Networks

Introduction

EIGRP and RIP are common Interior Gateway Protocols (IGP) used for routing in IP networks. Both protocols are advanced distance vector-based routing protocols. EIGRP is Cisco-proprietary.

In this chapter, we will perform:

- Routing protocol selection,
- Convergence analysis on link failures, and
- Network merge, route maps, and multiple EIGRP processes.

We will conclude that:

- EIGRP and RIP networks can be analyzed in NETWARS in terms of protocol overhead, convergence time, and route selection, and
- Complex operations such as merging two corporate networks running RIP or EIGRP can be modeled in NETWARS.

Note—The following examples were presented at OPNETWORK 2004 in Session 1312, Planning and Analyzing EIGRP and RIP Networks. If you do not have access to the files that these procedures use, you can still follow along using the sample screens provided in this user's guide.

There is a set of models with all of the procedural steps completed. You may use these for reference, if you do not wish to do the exercises yourself. These files end in the suffix "_ref" (reference).

Selecting a Routing Protocol

In the following example, a car manufacturing company wants to run a distance vector routing protocol over their network. They are considering RIP and EIGRP as possible candidates. Criteria to consider for selecting a routing protocol include: path selection; time taken for routing convergence; and protocol overhead.

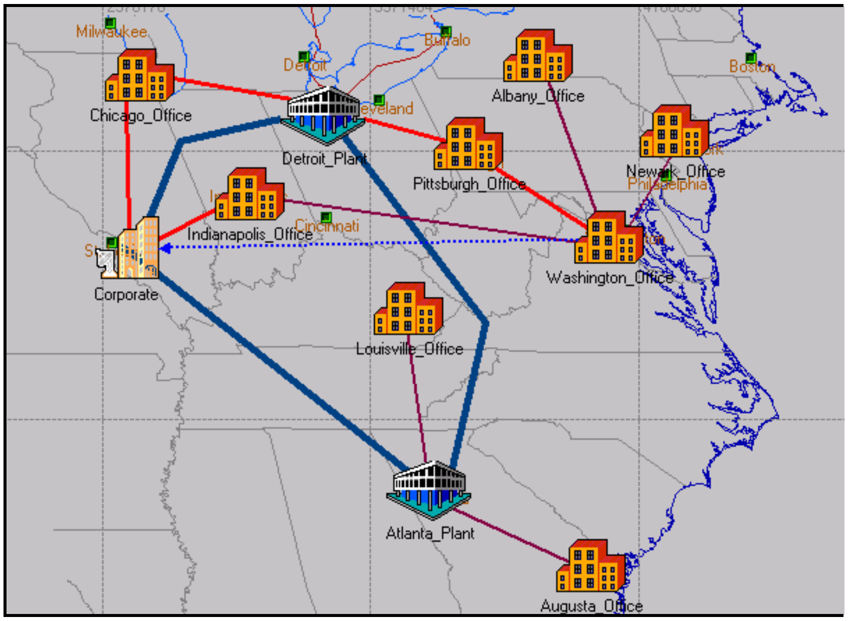

Figure 3-1 Network Model

In the following section, we will compare the difference in paths between two nodes with RIP and EIGRP. We'll view the routing convergence statistics (at start-up and in case of failure), and then we'll compare the routing protocol overhead.

Analyzing the RIP Routing Protocol

Procedure 3-1 Enable RIP on the Network Model

1 Launch NETWARS, if not already opened.

2 From the System Editor's **File** menu, choose **Open Editor**.

3 From the Open Editor drop-down menu, choose **Scenario Builder**, and then click **OK**.

4 Select **File > Open Project**. The Open Project dialog box displays.

5 In the Open Project dialog box, select the project file named `1312_lab1`, and then click **Open**.

Note—If you do not have access to this file, simply view the screens provided in this user's guide to follow along with the procedure.

6 Enable RIP on all interfaces in the network.

6.1 Select **Protocols > IP > Routing >Configure Routing Protocols...**.

6.2 Place a checkmark next to RIP. Notice the RIP selection is applied to all interfaces in the network.

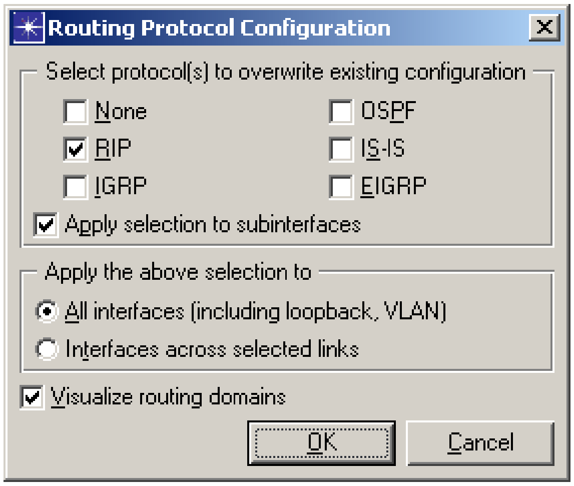

Routing Protocol Configuration

Select protocol(s) to overwrite existing configuration

□ None □ OSPF

☑ RIP □ IS-IS

□ IGRP □ EIGRP

☑ Apply selection to subinterfaces

Apply the above selection to

● All interfaces (including loopback, VLAN)

○ Interfaces across selected links

☑ Visualize routing domains

OK Cancel

Figure 3-2 Routing Protocol Configuration dialog box

6.3 Click **OK**.

7 Visually check the selection of RIP as routing protocol on all interfaces.

7.1 Verify that all the links in the network are annotated with “R”, showing that the interface on either side of the link runs RIP processes.

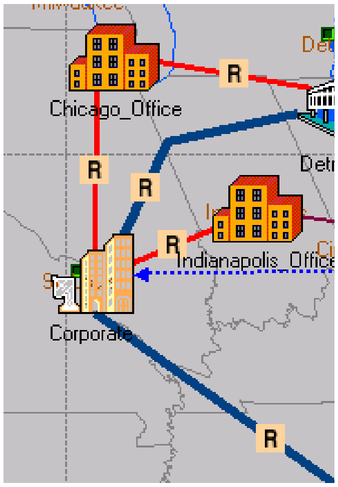

Figure 3-3 Links Showing RIP

7.2 To clear routing domain visualization, select **View > Visualize Protocol Configuration > Clear Visualization**.

8 Collect network-level convergence and protocol overhead statistics.

8.1 Right-click in the scenario space (away from any objects); select **Choose Individual DES Statistics**.

8.2 Expand the RIP category under Global Statistics by clicking on the + button next to it; select all the statistics under the RIP category, as shown below.

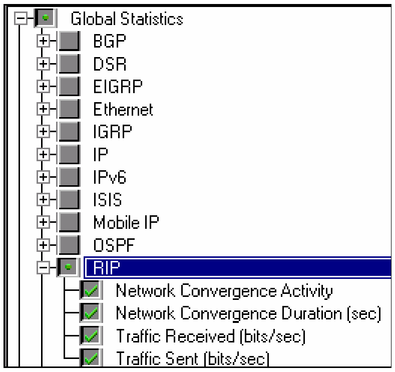

- Global Statistics
  - BGP
  - DSR
  - EIGRP
  - Ethernet
  - IGRP
  - IP
  - IPv6
  - ISIS
  - Mobile IP
  - OSPF
  - RIP
    - ☑ Network Convergence Activity
    - ☑ Network Convergence Duration (sec)
    - ☑ Traffic Received (bits/sec)
    - ☑ Traffic Sent (bits/sec)

Figure 3-4 Selecting All RIP Statistics

8.3 Click **OK** to confirm the statistics’ selection.

9 Run the simulation.

9.1 Click the **Configure/Run Simulation** toolbar button.

9.2 Click **Run** to start the simulation.

9.3 Close the Simulation Sequence dialog box after the simulation runs.

10 View the simulation results.

10.1 Right-click in the scenario space (away from any objects); select **View Results**.

10.2 In the RIP category, select **Network Convergence Activity** and **Network Convergence Duration**.

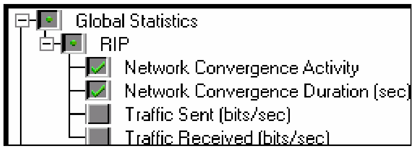

Global Statistics
RIP
☑ Network Convergence Activity
☑ Network Convergence Duration (sec)
□ Traffic Sent (bits/sec)
□ Traffic Received (bits/sec)

Figure 3-5 Selecting RIP Network Convergence Statistics

10.3 Click **Show** to display the graphs.

10.4 De-select **Network Convergence Activity** and **Network Convergence Duration**; then select **Traffic Sent** and **Traffic Received**.

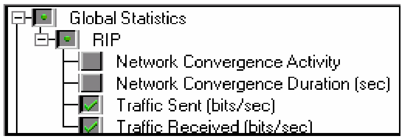

Global Statistics
RIP
□ Network Convergence Activity
□ Network Convergence Duration (sec)
☑ Traffic Sent (bits/sec)
☑ Traffic Received (bits/sec)

Figure 3-6 Selecting RIP Traffic Statistics

10.5 Click **Show** to display the graphs.

10.6 Close the View Results dialog box.

End of Procedure 3-1

Interpreting RIP Results

Procedure 3-2 Examine RIP on the Network Model

1 Examine RIP Network Convergence statistics.

The following panel displays the collected RIP Network Convergence statistics.

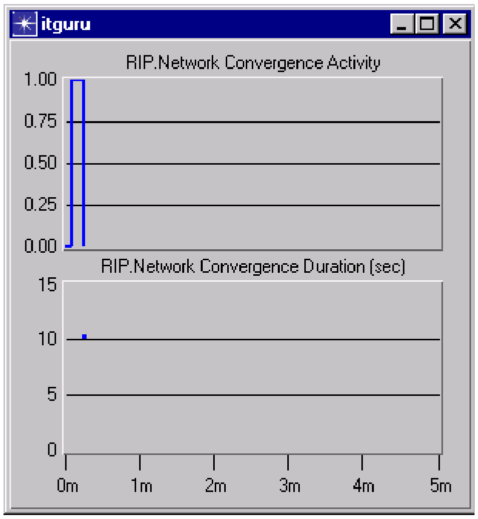

Figure 3-7 RIP Network Convergence Statistics

1.1 To zoom in the significant area of the graph: right-click on the edge of margin of the graph panel (but not inside the graph itself); choose **Edit Panel Properties** and edit the Horizontal Max to read "20s" (i.e. only the range between 0 and 20 seconds is relevant in this statistic); click **OK**.

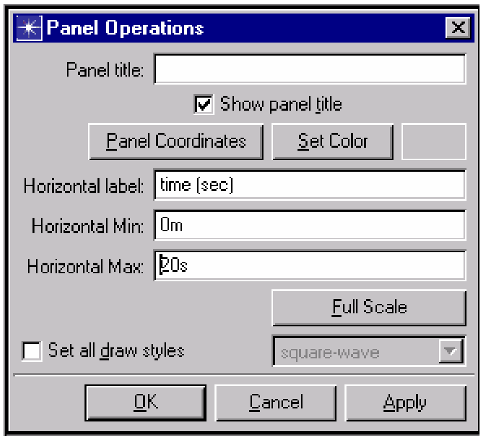
Panel Operations

Panel title:

☑ Show panel title

Horizontal label: time (sec)

Horizontal Min: 0m

Horizontal Max: 20s

□ Set all draw styles square-wave

Figure 3-8 Panel Operations dialog box

1.2 Notice in the resulting graphs that convergence of the RIP network starts at 5 seconds and finishes at little after 15 seconds, yielding a 10 second convergence duration. The two graphs display the same information from different points of view. The **RIP.Network Convergence Activity** graph plots the value of 1 for the time interval during which convergence activity is detected in the network, and the value of 0 for time intervals when there is no convergence activity. The **RIP.Network Convergence Duration** writes a point for each convergence activity interval – its x-value corresponds to the time convergence was achieved, while its y-value corresponds to the duration of the convergence interval.

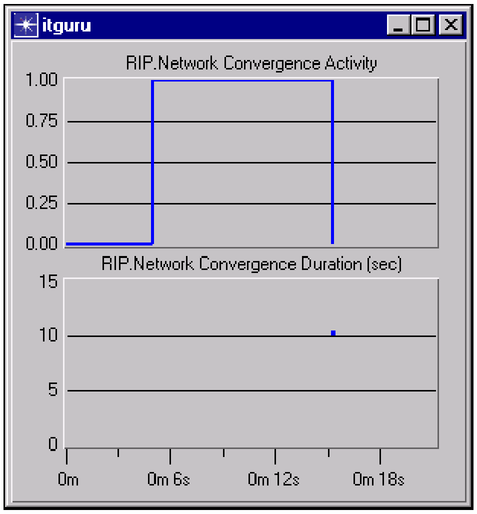

Figure 3-9 RIP Network Convergence Statistics

2 Examine RIP protocol overhead.

The following panel displays the collected RIP protocol overhead statistics.

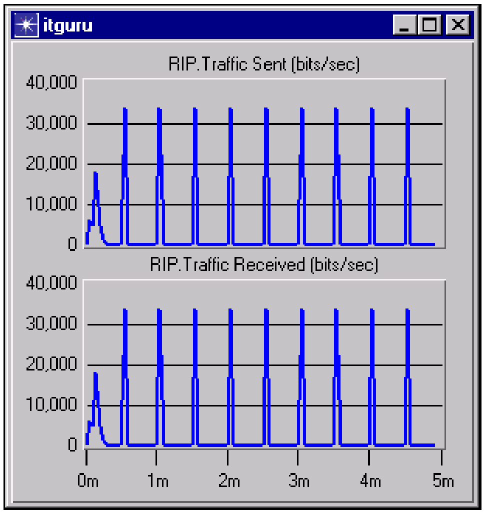

Figure 3-10 RIP Traffic Statistics

2.1 Notice that the **RIP.Traffic Sent** and **RIP.Traffic Received** graphs record the total protocol-generated traffic in the whole network. Notice a 30-seconds periodicity in the protocol-generated traffic, corresponding to the 30-second route update timer in RIP. You can infer that all the RIP processes in the network start at the same time, since their 30-second timers are synchronized to send updates at the same time, creating the peaks in the graphs followed by troughs of 0 bits/sec. Indeed all the RIP processes are configured to start at 5 seconds into the simulation. Had we staggered their start times, the convergence interval would likely be longer.

3 Examine the route chosen for the traffic flow between **Washington_Office** and **Corporate**.

3.1 Select **Protocols > IP > Demands >Display Routes for Configured Demands...**.

3.2 Expand the source **Washington_Office** by clicking on the (+) button, until you can see the route **Washington_Office** → **Corporate**, as shown below; click on the **Washington_Office** → **Corporate** route name.

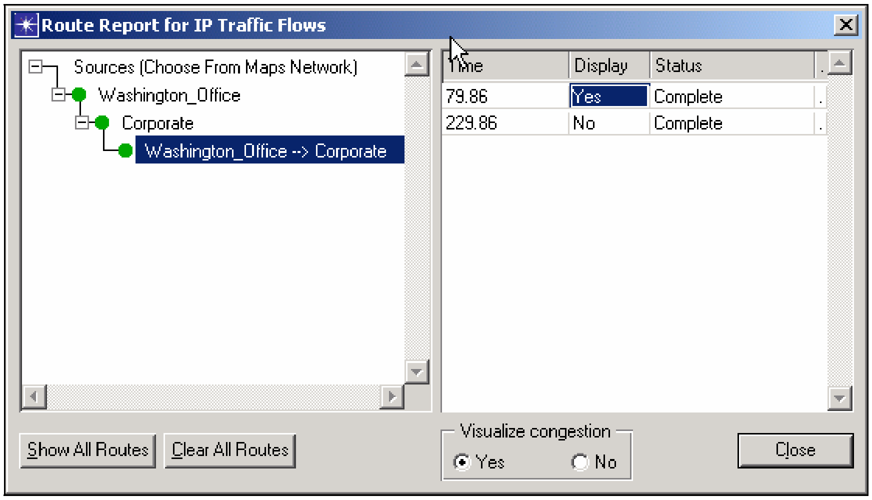

Route Report for IP Traffic Flows

- Sources (Choose From Maps Network)
  - Washington_Office
    - Corporate
      - Washington_Office --> Corporate

Time	Display	Status
79.86	Yes	Complete
229.86	No	Complete

Show All Routes | Clear All Routes

Visualize congestion: ● Yes ○ No

Close

Figure 3-11 Route Report for IP Traffic Flows dialog box

3.3 You should see two snapshots of the route at two different times in the right half of the panel; click on the Display field of any of the 2 snapshots to change the Display field to **Yes**, as shown above.

3.4 Look at the network to see the visual representation of the path the traffic flow was routed on: it was routed via the **Indianapolis_Office** node.

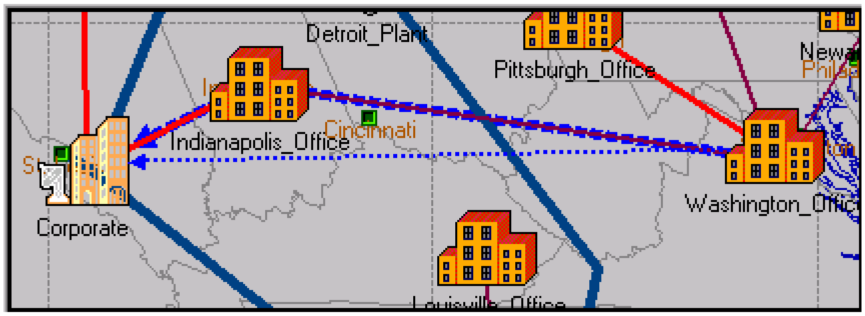

Figure 3-12 Traffic Routed via Indianapolis_Office

3.5 To clear the visual representation of the path, click **Clear All Routes** in the bottom left side of the Route Report for IP Traffic Flows dialog box.

3.6 Click **Close** to dismiss the Route Report for IP Traffic Flows dialog box.

3.7 Notice that the path **Washington_Office** → **Indianapolis_Office** → **Corporate** has the least number of hops. It uses the link **Washington_Office** → **Indianapolis_Office**, which has a rate of DS-1/1.544 Mbps (right-click on the link, choose **Edit Attributes > View Link Description** to confirm).

3.8 Notice there is another path between **Washington_Office** and **Corporate** which uses links of DS-3/44.736 Mbps. However, RIP cannot select a path based on available bandwidth. Instead, RIP always selects the path with the least number of hops.

The IT department of the car manufacturing company would like the path selection to be based on available bandwidth. Hence, they decide to try running EIGRP on the network.

End of Procedure 3-2

Analyzing the EIGRP Routing Protocol

Procedure 3-3 Enable EIGRP on the Network Model

1 Select **Scenario > Switch to Scenario**, and choose scenario *EIGRP*. A new scenario displays with the same network, already configured to run EIGRP on all interfaces.

2 Visually check that EIGRP is configured on all interfaces.

2.1 Select **View > Visualize Protocol Configuration > IP Routing Domains**.

2.2 Notice that all links are annotated with an “E”. Such a link annotation means that EIGRP is active on the interfaces at both ends of the link.

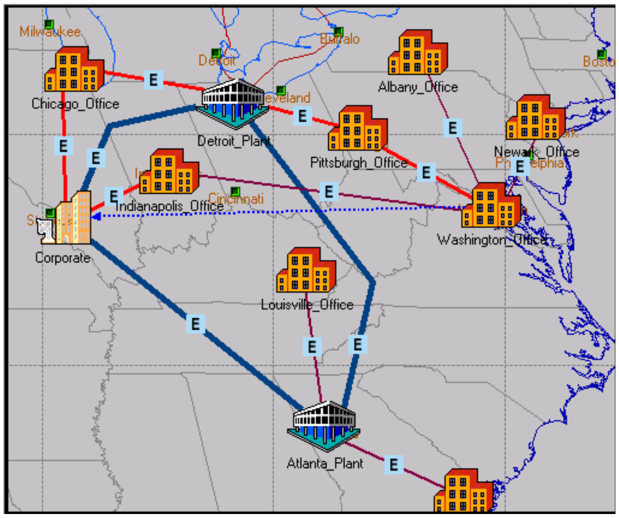

Figure 3-13 Links Showing EIGRP

3 Collect network-level convergence and protocol overhead statistics.

3.1 Right-click in the scenario space (away from any objects); select **Choose Individual DES Statistics**.

3.2 Expand the EIGRP category under Global Statistics by clicking on the + button next to it; select all the statistics under the EIGRP category, as shown below.

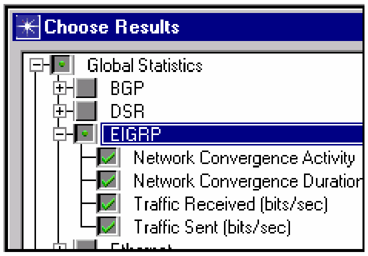
Choose Results
Global Statistics
BGP
DSR
EIGRP
Network Convergence Activity
Network Convergence Duration
Traffic Received (bits/sec)
Traffic Sent (bits/sec)

Figure 3-14 Selecting All EIGRP Statistics

3.3 Click **OK** to confirm the statistics' selection.

4 Run the simulation.

4.1 Click the **Configure/Run Simulation** toolbar button.

4.2 Click **Run** to start the simulation.

4.3 Close the Simulation Sequence dialog box after the simulation runs.

5 View the simulation results.

5.1 Panel template have already been saved with the scenario; to retrieve the saved panels, click the **Show/Hide Graphs** button.

5.2 Load the panels with latest results by selecting **DES > Panel Operations > Panel Templates > Load with Latest Results**.

You should now be able to see the latest results in the panels.

End of Procedure 3-3

Interpreting EIGRP Results

Procedure 3-4 Examine EIGRP on the Network Model

1 Compare EIGRP versus RIP convergence statistics.

The following panel displays the collected EIGRP Network Convergence statistics.

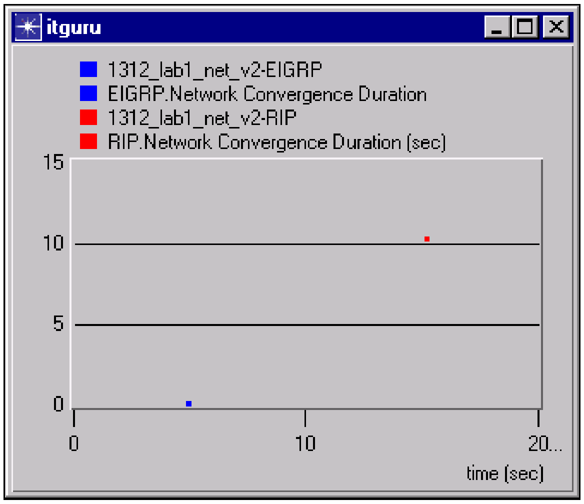

Figure 3-15 EIGRP Network Convergence Statistics

1.1 To see what the x-value and y-value are for a point recorded in the statistic, simply place the mouse tip close to the point of interest – a box will appear showing the coordinates of the nearest recorded point.

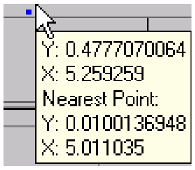

Y: 0.4777070064
X: 5.259259
Nearest Point:
Y: 0.0100136948
X: 5.011035

Figure 3-16 Coordinates Tip

Notice that the EIGRP network takes approx. 10^{-2} seconds to converge, while the same network running RIP took more than 10 seconds to converge. The results seem puzzling: where would a 3 orders of magnitude difference come from? It turns out this difference is due to a RIP requirement designed to limit the frequency of triggered updates. RFC 1058 states: “After a triggered update is sent, a timer should be set for a random time between 1 and 5 seconds. If other changes occur before the timer expires, a single update is triggered when the timer expires.” This mechanism meant to prevent an excessive load on the network due to triggered updates, can also dramatically extend the convergence time.

2 Compare EIGRP versus RIP protocol overhead.

The following panel displays the EIGRP versus RIP protocol overhead statistics.

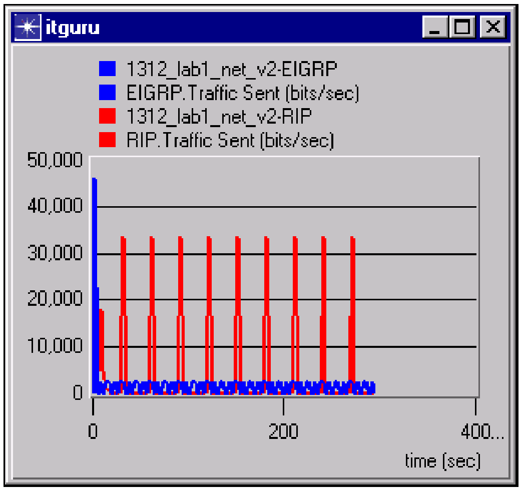

Figure 3-17 EIGRP versus RIP Protocol Overhead Statistics

Notice that RIP traffic peaks much higher than the EIGRP traffic in 30-second intervals. This corresponds to the fact that RIP re-sends the whole topology information every 30 seconds, while EIGRP sends the information once, then relies on short hello messages to keep the information up to date.

End of Procedure 3-4

Analyzing RIP and EIGRP in a Failure-Recovery Scenario

Two additional scenarios have been created for a study comparing the behavior of RIP and EIGRP in case of a link failure, followed by recovery.

Procedure 3-5 Set Up RIP and EIGRP Failure-Recovery Scenario

1 Inspect settings in the RIP_failure scenario.

1.1 Choose **Scenario > Switch to Scenario > RIP_failure**.

1.2 Note the network in RIP_failure is the same as the one in RIP, except for the addition of a failure-recovery definition object and the replacement of the traffic flow between **Washington_Office** and **Corporate** with a flow between **Indianapolis_Office** and **Corporate**. The link between **Indianapolis_Office** and **Corporate** will fail and subsequently recover – we placed a traffic flow across this link in order to monitor how this traffic flow will be routed during the link failure.

1.3 Right-click on the failure-recovery object and choose **Edit Attributes**.

1.4 To examine the failure settings of the **Indianapolis_Office** → **Corporate link**, double-click in the Value column corresponding to the Link Failure/Recovery Specification attribute.

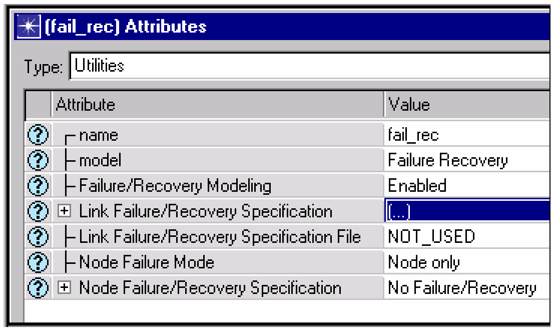

(fail_rec) Attributes

Type: Utilities

Attribute	Value
name	fail_rec
model	Failure Recovery
Failure/Recovery Modeling	Enabled
⊞ Link Failure/Recovery Specification	(...)
Link Failure/Recovery Specification File	NOT_USED
Node Failure Mode	Node only
⊞ Node Failure/Recovery Specification	No Failure/Recovery

Figure 3-18 Link Failure/Recovery Specification Attribute

1.5 Notice that the **Indianapolis_Office↔Corporate** link fails at 100 seconds and recovers at 200.

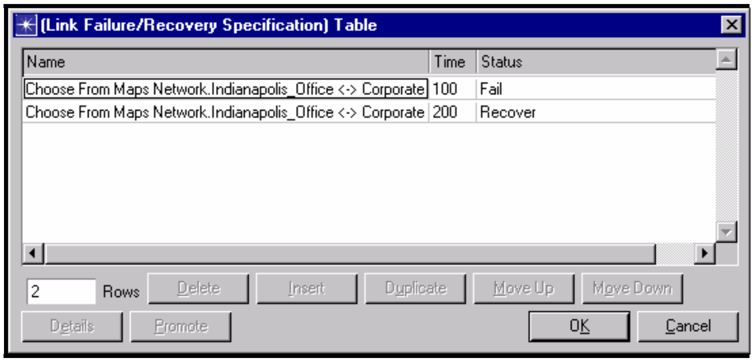

(Link Failure/Recovery Specification) Table

Name	Time	Status
Choose From Maps Network.Indianapolis_Office <-> Corporate	100	Fail
Choose From Maps Network.Indianapolis_Office <-> Corporate	200	Recover

2 Rows

Figure 3-19 Link Failure/Recovery Specification Attribute Table

2 Inspect settings in the “EIGRP_failure” scenario.

2.1 Choose **Scenario > Switch to Scenario > EIGRP_failure**.

2.2 Note that the network in EIGRP_failure is the same as the one in RIP_failure scenario that you have just examined except that EIGRP is configured to run on all interfaces instead of RIP. As in RIP_failure, in this scenario, the link between **Indianapolis_Office** and **Corporate** will fail at 100 seconds and will recover at 200 seconds.

3 Compare results between RIP and EIGRP deployed in a failure scenario.

3.1 Panel templates for comparison statistics between RIP and EIGRP have already been saved in scenario EIGRP_failure. Bring up the panel templates by clicking the **Show/Hide Graphs** button.

3.2 To load the panels with the latest results, select **DES > Panel Operations > Panel Templates > Load with Latest Results**.

You should be able now to see the latest results in the panels.

End of Procedure 3-5

Comparing and Interpreting RIP and EIGRP Results

Procedure 3-6 Compare and Interpret RIP and EIGRP Results

1 Compare EIGRP versus RIP convergence statistics.

The following panel displays a comparison of RIP and EIGRP convergence times.

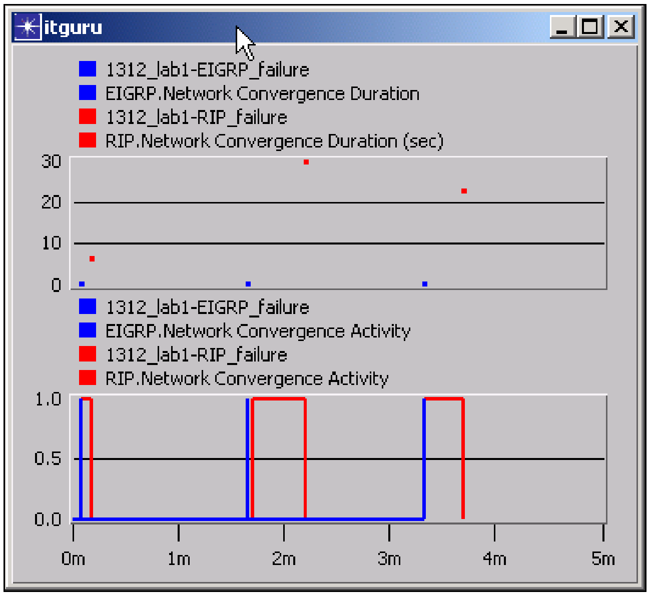
itguru
🟦 1312_lab1-EIGRP_failure
🟦 EIGRP.Network Convergence Duration
🟥 1312_lab1-RIP_failure
🟥 RIP.Network Convergence Duration (sec)
30
20
10
0
🟦 1312_lab1-EIGRP_failure
🟦 EIGRP.Network Convergence Activity
🟥 1312_lab1-RIP_failure
🟥 RIP.Network Convergence Activity
1.0
0.5
0.0
0m
1m
2m
3m
4m
5m

Figure 3-20 RIP and EIGRP Network Convergence Statistics

1.1 The two graphs show the same data, albeit with different emphases. The Network Convergence Duration shows the time it took the network to achieve convergence as (x, y) points. There are 3 pairs of points corresponding to 3 convergence intervals: the first one happens after initialization, the second one happens after the configured link failure at 100 seconds, the third one after link recovery at 200 sec. The Network Convergence Activity statistics shows a value of 1 on the y-axis for as long as there is convergence activity in the network.

1.2 Notice that following a disturbance (such a link failure, or a link recovery), convergence is achieved in approx. 10^{-2} seconds in the EIGRP network, versus anywhere between 10 and 20 seconds in the RIP network.

2 Compare EIGRP and RIP routing protocol overhead.

The following graph displays a comparison of routing protocol overhead in the same failure-recovery network for RIP and EIGRP.

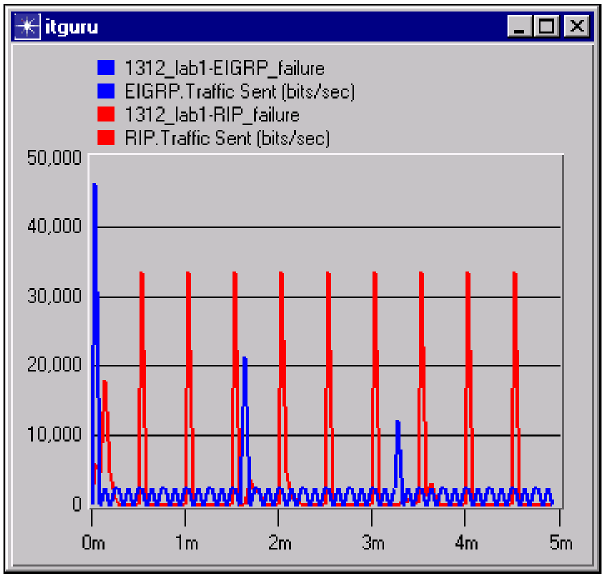

Figure 3-21 EIGRP and RIP Routing Protocol Overhead

2.1 Notice that EIGRP routing protocol overhead increases visibly around the link failure at 100 seconds, and again it increases around the link recovery at 200 seconds, but stays at a low value at all other times. On the other hand, RIP overhead is periodic, and it is - relative to the activity peak at every 30 seconds – less influenced by the failure-recovery events.

3 Compare EIGRP and RIP routing of traffic flow between **Indianapolis_Office** and **Corporate** before, during and after link failure.

Repeat the steps below for EIGRP_failure and RIP_failure scenarios. Start with EIGRP_failure scenario.

3.1 In case the link **Indianapolis_Office↔Corporate** fails, notice that the traffic flow between **Indianapolis_Office** and **Corporate** will have to be routed via **Washington_Office**. From **Washington_Office** to **Corporate**, notice there is one routing solution that minimizes the number of hops at the same time as it maximizes bandwidth via **Pittsburg_Office** and **Detroit_Plant**. Consequently, we expect both RIP and EIGRP to choose this path during the failure of link **Indianapolis_Office↔Corporate**.

3.2 Choose **Protocols > IP > Demands > Display Routes for Configured Demands…**.

3.3 On the left side of the dialog box, click on the + button next to **Indianapolis_Office** to expand the hierarchy as shown below; click on the traffic flow name **Indianapolis_Office→Corporate** to see in the right panel snapshots of the route taken by the traffic flow at different times.

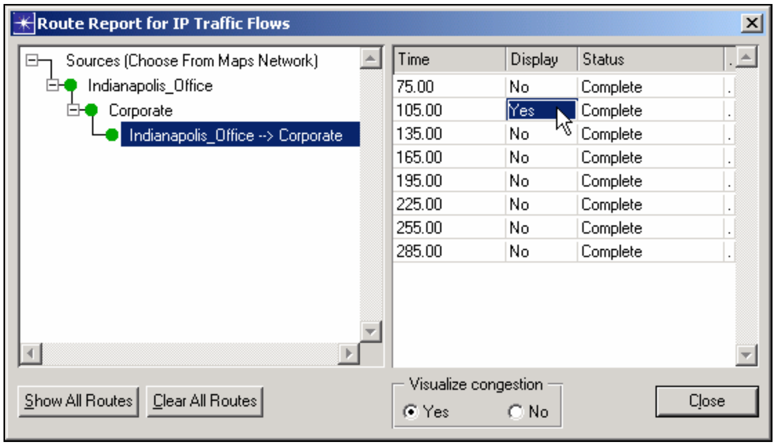

Route Report for IP Traffic Flows

Sources (Choose From Maps Network)
- Indianapolis_Office
  - Corporate
    - Indianapolis_Office --> Corporate

Time	Display	Status
75.00	No	Complete
105.00	Yes	Complete
135.00	No	Complete
165.00	No	Complete
195.00	No	Complete
225.00	No	Complete
255.00	No	Complete
285.00	No	Complete

Visualize congestion: ● Yes ○ No

Figure 3-22 Route Report for IP Traffic Flows dialog box

3.4 Click on the Display column, in the row corresponding to a given time, to change the display from **No** to **Yes**. Look at the network to visualize the route taken by the traffic flow at that time. For instance, below is a snapshot of the route at 135 seconds – since it falls between 100 and 200 seconds, the link **Corporate** ↔ **Indianapolis_Office** is failed and the traffic flow is routed as discussed above.

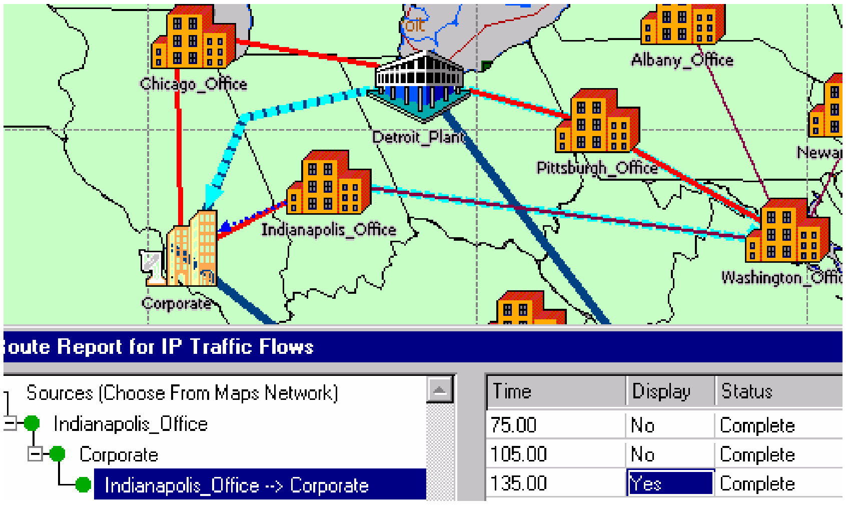

oute Report for IP Traffic Flows

Sources (Choose From Maps Network)
- Indianapolis_Office
  - Corporate
    - Indianapolis_Office --> Corporate

Time	Display	Status
75.00	No	Complete
105.00	No	Complete
135.00	Yes	Complete

Figure 3-23 Route Snapshot at 135 Seconds

3.5 Clicking again in the Display column will change the value of **Yes** to **No** and the route will no longer be displayed; experiment with viewing snapshots of the route at different times, bearing in mind that between 100 and 200 seconds the link **Indianapolis_Office↔Corporate** is not available for routing.

3.6 To verify the routing of the same traffic flow in the RIP scenario, choose **Scenario > Switch To Scenario > RIP_failure**.

3.7 Repeat steps above in the RIP_failure scenario. Notice that RIP will route the demand during failure in the same way as EIGRP, since it so happens that a path satisfied both the least number of hops requirement (for RIP) and the bandwidth requirement (for EIGRP). Also, note that the demand at 105 sec is unroutable for RIP since RIP has not yet converged at that time.

End of Procedure 3-6

Conclusion

In the previous section, we introduced the network of a car manufacturer company and compared the performance of RIP and EIGRP over the introduced network. Three aspects were compared: route choice between RIP and EIGRP, routing protocol overhead and routing protocol convergence.

From the point of view of route choice between RIP and EIGRP, we saw a situation where a traffic flow is routed differently for RIP and for EIGRP – furthermore, we concluded that RIP is limited in that it only considers the least number of hops as its optimality criterion. In the failure recovery scenario, we also saw that a traffic flow is routed through the same path in both RIP and EIGRP.

From the routing protocol overhead point of view, we saw that RIP uses more bandwidth than EIGRP, even in the case of a failure recovery scenario.

From the point of view of routing protocol convergence time, we saw that there is a 3-order of magnitude difference between the time EIGRP takes to converge and the time RIP takes to converge over the same network, over 3 different situations: network start-up, link failure and link recovery. More exactly, EIGRP takes approx 10^{-2} seconds to convergence, when RIP takes 10 seconds or more to converge. The root cause for this difference lies in a RIP RFC specification, according to which a timer anywhere between 1 and 5 seconds limits the frequency of triggered updates.

Merging Two Networks

In the previous sections of this chapter, we saw a car manufacturing company choose a routing protocol to run over their corporate network. The candidates were RIP and EIGRP. From the point of view of traffic engineering, routing convergence times and routing protocol overhead, EIGRP appeared to be more beneficial. In this section, the car manufacturing company who is using EIGRP in their corporate network is merging with a tire manufacturing company, Tirex, who is using RIP to route traffic.

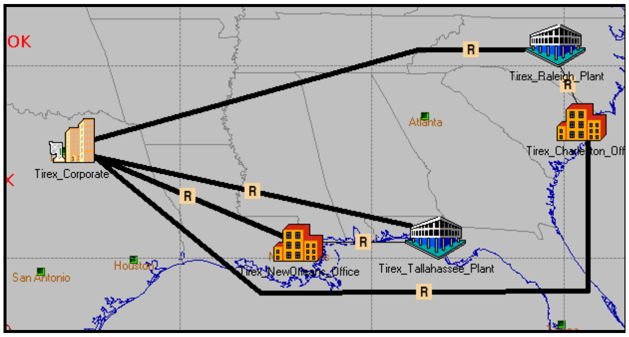

Figure 3-24 Links Showing RIP

In the following procedure, we will merge the two networks (configure redistribution on border routers), and then use EIGRP in the whole network.

First, we will examine the Tirex network, and then we will configure the two networks for merger. We will use route maps to prevent routes from being redistributed back to the original network. Then we will migrate the Tirex network from RIP to EIGRP, and use two EIGRP processes.

Examining the Tirex Network

Procedure 3-7 Examine the Tirex Network

1. Launch NETWARS, if not already opened.
2. From the System Editor’s **File** menu, choose **Open Editor**.
3. From the Open Editor drop-down menu, choose **Scenario Builder**, and then click **OK**.
4. Select **File > Open Project**. The Open Project dialog box displays.

5 In the Open Project dialog box, select the project file named `1312_lab2`, and then click **Open**.

Note—If you do not have access to this file, simply view the screens provided in this user's guide to follow along with the procedure.

6 Visualize routing protocol configuration.

6.1 Select **View > Visualize Protocol Configuration > IP Routing Domains**.

All the links in the network appear annotated with the letter "R", meaning the interface at either end of each link is running RIP routing protocol.

7 Press **Ctrl+Shift+C** to clear the routing protocol visualization.

End of Procedure 3-7

Configuring the Networks for Merger

Procedure 3-8 Configure the Networks for Merger

1 Select **Scenario > Switch to Scenario**, and choose scenario "Merged_Network_Redist".

Notice that the topologies of the two corporate networks have been connected via two links: a T-3 link running between **Corporate** and **Tirex_Corporate** and a DS-1 link between **Washington_Office** and **Tirex_Raleigh_Plant**.

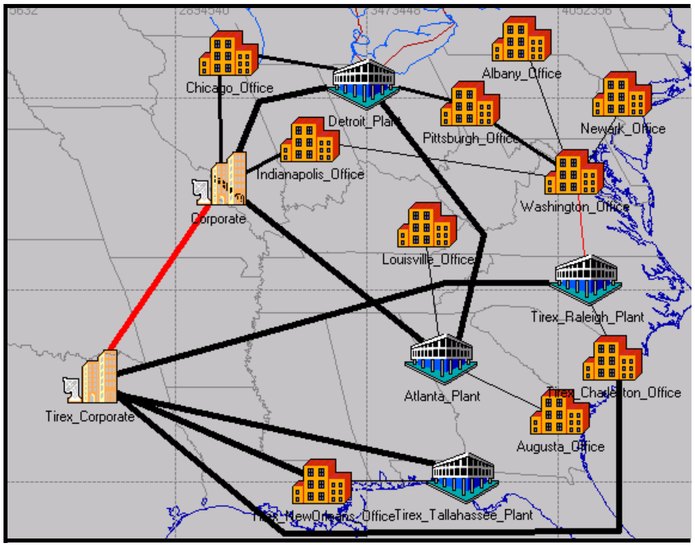

Figure 3-25 Two Networks Connected

2 Configure EIGRP on the interfaces between the two networks.

2.1 Select the newly added links **Corporate↔Tirex_Corporate** and **Washington_Office↔Tirex_Raleigh_Plant** (press **Shift+down** after clicking on one to select the other as well).

2.2 Select **Protocols > IP > Routing >Configure Routing Protocols...**.

2.3 Select EIGRP by placing a checkmark next to the protocol name; un-check RIP (also make sure no other protocols are checked); select **Interfaces across selected links**; click **OK**.

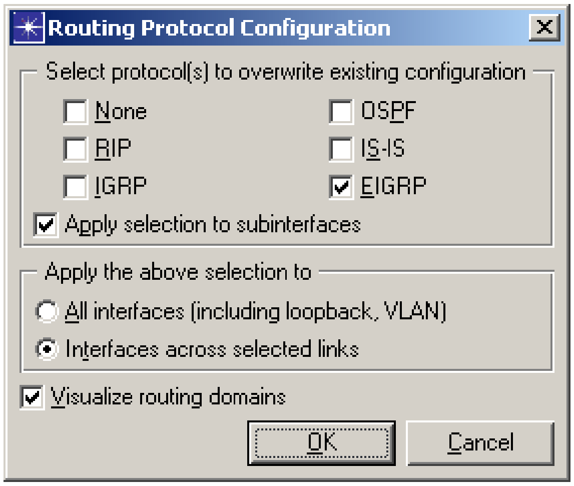

Routing Protocol Configuration

Select protocol(s) to overwrite existing configuration

☐ None ☐ OSPF

☐ RIP ☐ IS-IS

☐ IGRP ☑ EIGRP

☑ Apply selection to subinterfaces

Apply the above selection to

○ All interfaces (including loopback, VLAN)

● Interfaces across selected links

☑ Visualize routing domains

Figure 3-26 Routing Protocol Configuration dialog box

2.4 In the ensuing routing domain visualization, notice the two selected links have an ‘E’ annotation, confirming that EIGRP is running across these links.

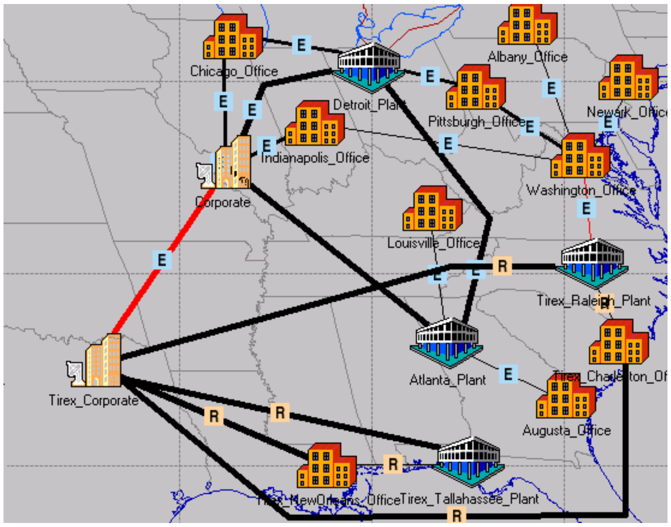

Figure 3-27 Selecting EIGRP Links

2.5 Clear the routing protocol visualization by pressing **Ctrl+Shift+C**.

3 Configure redistribution between RIP and EIGRP.

3.1 Select the nodes **Tirex_Raleigh_Plant** and **Tirex_Corporate** by holding down the **Shift** key while clicking on the two nodes.

3.2 Redistribution into EIGRP:

- Select **Protocols > EIGRP > Configure Route Redistribution...**.
- In the RIP row, click in the Status column to change value to **Enabled**; make sure the change applies to selected routers, by clicking the appropriate radio button.

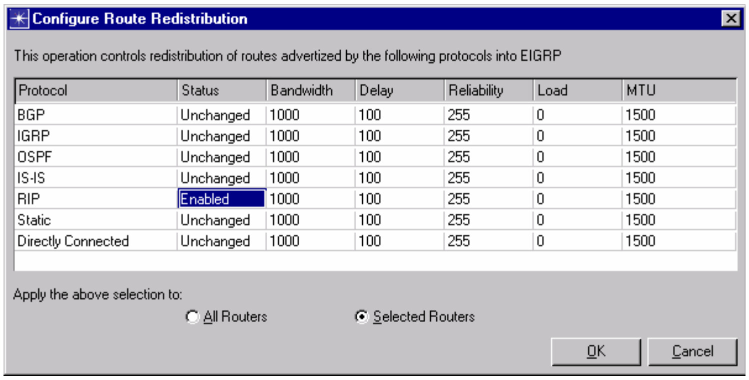

This operation controls redistribution of routes advertized by the following protocols into EIGRP

Protocol	Status	Bandwidth	Delay	Reliability	Load	MTU
BGP	Unchanged	1000	100	255	0	1500
IGRP	Unchanged	1000	100	255	0	1500
OSPF	Unchanged	1000	100	255	0	1500
IS-IS	Unchanged	1000	100	255	0	1500
RIP	Enabled	1000	100	255	0	1500
Static	Unchanged	1000	100	255	0	1500
Directly Connected	Unchanged	1000	100	255	0	1500

Apply the above selection to:

○ All Routers ● Selected Routers

Figure 3-28 RIP Status in the Configure Route Redistribution dialog box

- Click **OK** to dismiss the Route Redistribution window.

3.3 Redistribution into RIP:

- Select **Protocols > RIP > Configure Route Redistribution…**
- In the EIGRP row, click in the Status column to change value to **Enabled**; make sure the change applies to selected routers, by clicking the appropriate radio button.

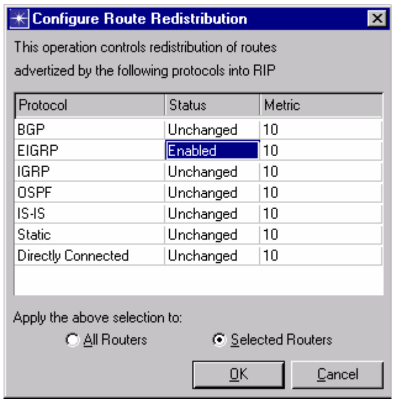

This operation controls redistribution of routes advertized by the following protocols into RIP

Protocol	Status	Metric
BGP	Unchanged	10
EIGRP	Enabled	10
IGRP	Unchanged	10
OSPF	Unchanged	10
IS-IS	Unchanged	10
Static	Unchanged	10
Directly Connected	Unchanged	10

Apply the above selection to:

○ All Routers ● Selected Routers

Figure 3-29 EIGRP Status in the Configure Route Redistribution dialog box

- Click **OK** to dismiss the Route Redistribution dialog box.

4 Configure route maps to filter route updates.

Since we have configured redistribution between the two networks, we have to make sure that routes from the Tirex network, once distributed in the original network, are not advertised back into the Tirex network. Similarly, routes from the original network, once distributed in the Tirex network, should not be advertised back into the original network. To selectively filter route updates, we have defined route maps in the two routers involved in the redistribution process: **Tirex_Raleigh_Plant** and **Tirex_Corporate**.

4.1 To examine the route maps definition on **Tirex_Corporate**, right-click on the node and choose **Edit Attributes**; expand the IP group by clicking on the + sign; select **IP Routing Parameters > Route Map Configuration**.

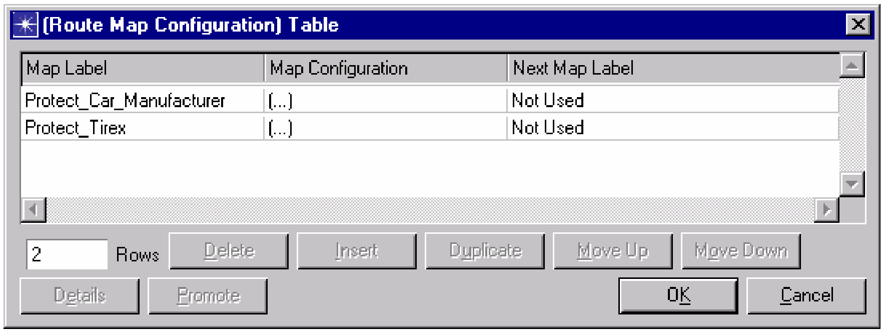

(Route Map Configuration) Table

Map Label	Map Configuration	Next Map Label
Protect_Car_Manufacturer	(...)	Not Used
Protect_Tirex	(...)	Not Used

2 Rows

Figure 3-30 Route Map Configuration Table for Tirex_Corporate node

4.2 Two route maps are defined on the **Tirex_Corporate** node: **Protect_Car_Manufacturer** and **Protect_Tirex**. To examine the definition of **Protect_Car_Manufacturer**, double-click in the Map Configuration field.

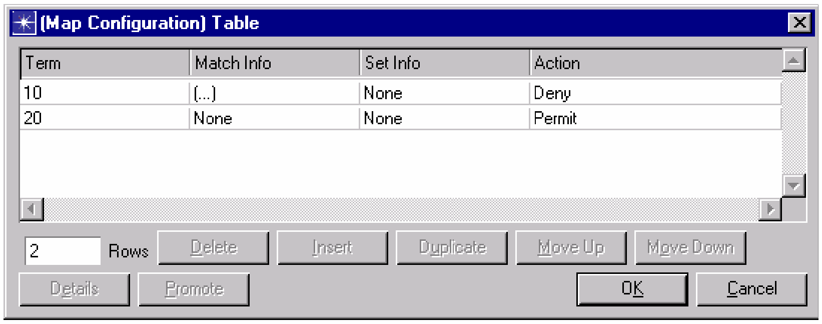

(Map Configuration) Table

Term	Match Info	Set Info	Action
10	(...)	None	Deny
20	None	None	Permit

2 Rows

Figure 3-31 Map Configuration Table dialog box for Protect_Car_Manufacturer Route

4.3 Create a Route Map that Prevents Distribution into Car Manufacturer.

- The **Protect_Car_Manufacturer** route map has two rows in its definition: the second row contains the default action of allowing traffic through the interface on which this route map will be installed; the first row contains an action of denying access to traffic that matches the condition defined in the Match Info field. Double-click on Match Info in the first row in order to see the matching condition for traffic being denied access.

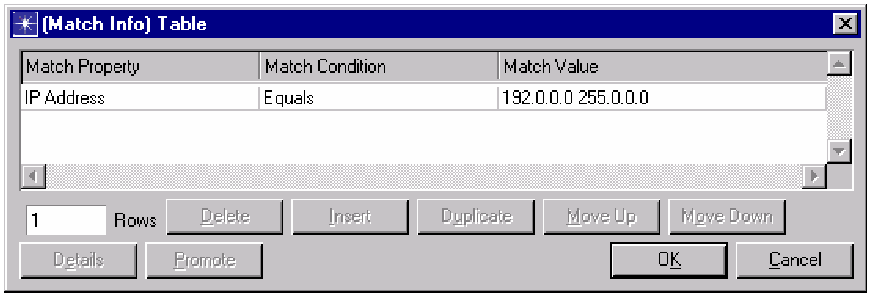

Match Property	Match Condition	Match Value
IP Address	Equals	192.0.0.0 255.0.0.0

Rows: 1

Figure 3-32 Match Info Table dialog box for Protect_Car_Manufacturer Route

- Notice that routing updated destined for the prefix 192.0.0.0/8 are denied access. This prefix corresponds to the original network of the car manufacturer. Consequently, **Protect_Car_Manufacturer** can be used to prevent routes to destinations inside the original network from coming into the original network from the Tirex network.
- Click **Cancel** to dismiss all attribute windows.

4.4 Create Route map that Prevents Distribution into Tirex.

- **Protect_Tirex** has been defined in a similar way, except that routes destined for the prefix 200.0.0.0/8 are denied access. This prefix belongs to the Tirex network. Hence **Route Map 2** can be used to prevent routes to destinations in the Tirex network from being advertised into the Tirex network from the original network.

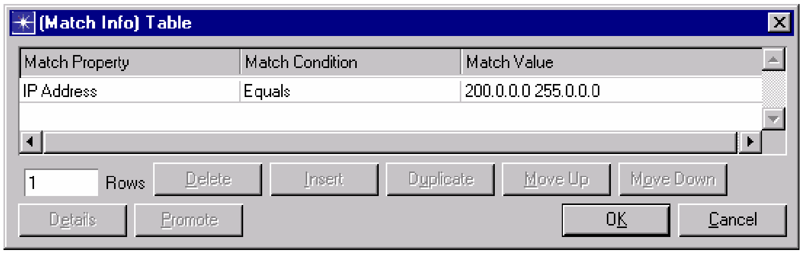

Match Property	Match Condition	Match Value
IP Address	Equals	200.0.0.0 255.0.0.0

Rows: 1

Figure 3-33 Match Info Table dialog box for Protect_Tirex Route

We have examined the definition of two route maps. The **Protect_Car_Manufacturer** route map is designed to protect the car manufacturer’s network from loops arising from redistribution, while the **Protect_Tirex** route map is designed to protect the Tirex Network for the same reason. In order for the route maps to be of any use, they need to be installed among the parameters that control redistribution between the RIP and EIGRP processes active on the redistribution-enabled nodes: **Tirex_Corporate** and **Tirex_Raleigh_Plant**.

4.5 Install Route Map for EIGRP Redistribution.

- To install **Protect_Car_Manufacturer** route map on the **Tirex_Corporate** node, right click on the node and choose **Edit Attributes**; expand the IP Routing Protocols category by clicking the + button; choose **EIGRP Parameters > Process Parameters > Process Parameters > Address Family Parameters > Redistribution > Routing Protocols > RIP**.

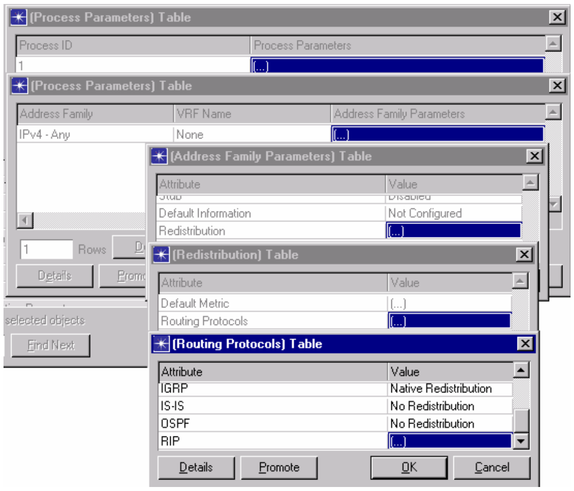

(Process Parameters) Table

Process ID	Process Parameters
1	(...)

(Process Parameters) Table

Address Family	VRF Name	Address Family Parameters
IPv4 - Any	None	(...)

(Address Family Parameters) Table

Attribute	Value
Default Information	Not Configured
Redistribution	(...)

(Redistribution) Table

Attribute	Value
Default Metric	(...)
Routing Protocols	(...)

(Routing Protocols) Table

Attribute	Value
IGRP	Native Redistribution
IS-IS	No Redistribution
OSPF	No Redistribution
RIP	(...)

Figure 3-34 Routing Protocols Table dialog box

- Click in value field of the Route Map attribute and select **Protect_Car_Manufacturer** from the drop down list.

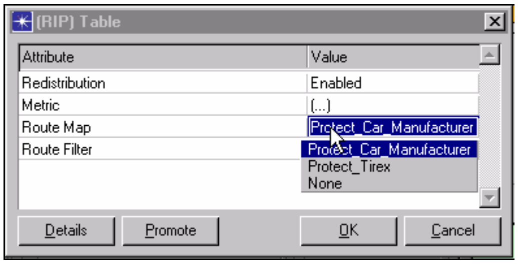

(RIP) Table

Attribute	Value
Redistribution	Enabled
Metric	(...)
Route Map	Protect_Car_Manufacturer
Route Filter	

Figure 3-35 RIP Table dialog box

- Click **OK** to dismiss all attribute windows.

4.6 Install Route Map for RIP Redistribution.

- To install **Protect_Tirex** on the **Tirex_Corporate** node, right-click on the node and choose **Edit Attributes**; expand the IP Routing Protocols category by clicking the (+) button; choose **RIP Parameters > Process Parameters > Process Parameters > Redistribution > Routing Protocols > EIGRP**.

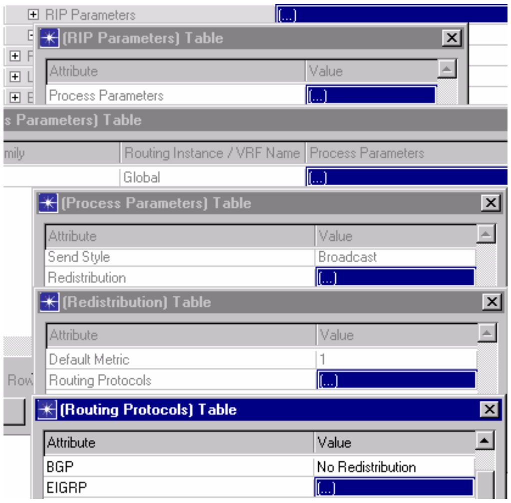

RIP Parameters (...)

(RIP Parameters) Table

Attribute	Value
Process Parameters	(...)

s Parameters) Table

mily	Routing Instance / VRF Name	Process Parameters
	Global	(...)

(Process Parameters) Table

Attribute	Value
Send Style	Broadcast
Redistribution	(...)

(Redistribution) Table

Attribute	Value
Default Metric	1
Routing Protocols	(...)

(Routing Protocols) Table

Attribute	Value
BGP	No Redistribution
EIGRP	(...)

Figure 3-36 Routing Protocols Table dialog box

- Click in the Route Map field to select **Protect_Tirex** from the drop down list.

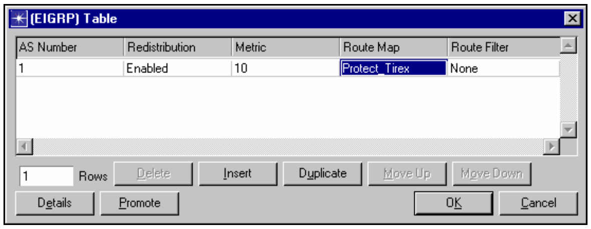

(EIGRP) Table

AS Number	Redistribution	Metric	Route Map	Route Filter
1	Enabled	10	Protect_Tirex	None

1 Rows

Figure 3-37 EIGRP Table dialog box

- Click **OK** to dismiss all attribute windows.
- The two route maps have been set in a similar way in the other route-redistribution node **Tirex_Raleigh_Plant**.

5 Import traffic flows from file.

5.1 Select **Traffic > Import Flows > From Spreadsheet**.

5.2 Choose the file `traffic`, located in directory `C:\op_models` then click **Open**.

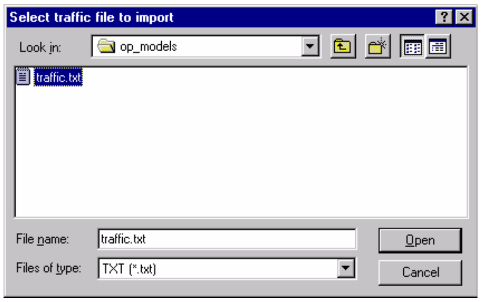

File name: traffic.txt

Files of type: TXT (*.txt)

Figure 3-38 Select Traffic File to Import dialog box

5.3 Click **OK**.

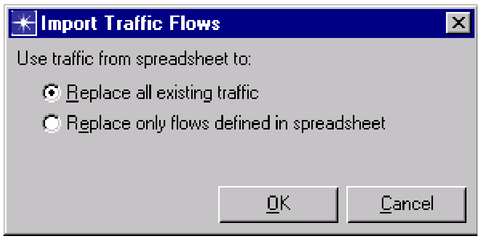

Use traffic from spreadsheet to:

● Replace all existing traffic

○ Replace only flows defined in spreadsheet

Figure 3-39 Import Traffic Flows dialog box

5.4 When the import finishes, import statistics are summarized. Click **Close** to dismiss them.

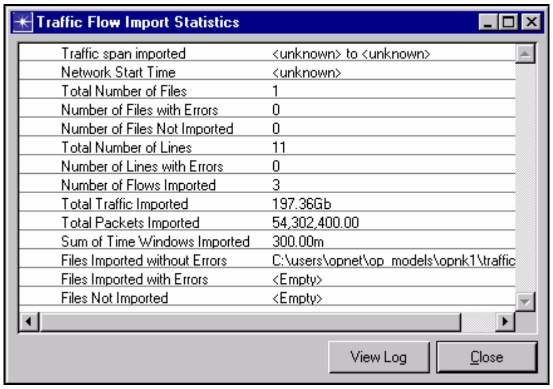

Traffic span imported	<unknown> to <unknown>
Network Start Time	<unknown>
Total Number of Files	1
Number of Files with Errors	0
Number of Files Not Imported	0
Total Number of Lines	11
Number of Lines with Errors	0
Number of Flows Imported	3
Total Traffic Imported	197.36Gb
Total Packets Imported	54,302,400.00
Sum of Time Windows Imported	300.00m
Files Imported without Errors	C:\users\opnet\op_models\opnk1\traffic
Files Imported with Errors	<Empty>
Files Not Imported	<Empty>

Figure 3-40 Traffic Flow Import Statistics dialog box

6 Run the simulation.

6.1 Click the **Configure/Run Simulation** toolbar button.

6.2 Click **Run** to start the simulation.

6.3 Close the Simulation Sequence dialog box after the simulation runs.

7 Color links by utilization.

7.1 To visualize the utilization of different links, select **View > Visualize Link Loads > Color by Link Load**; a window describing the feature pops up – click **OK**.

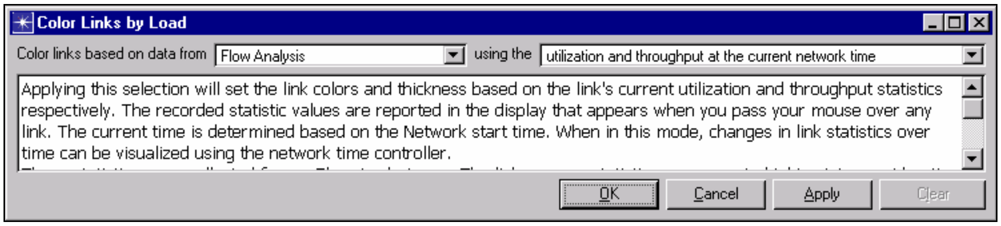

Color Links by Load

Color links based on data from: Flow Analysis — using the: utilization and throughput at the current network time

Applying this selection will set the link colors and thickness based on the link's current utilization and throughput statistics respectively. The recorded statistic values are reported in the display that appears when you pass your mouse over any link. The current time is determined based on the Network start time. When in this mode, changes in link statistics over time can be visualized using the network time controller.

Figure 3-41 Color Links by Load dialog box

7.2 The links of non-zero utilization in the network are color-coded as follows: green (links between 0-50% utilized), yellow (links between 50-75% utilized) and red (links more than 75% utilized). The links that are not utilized at all remain blue. The color-coding is independent in the two directions of the link.

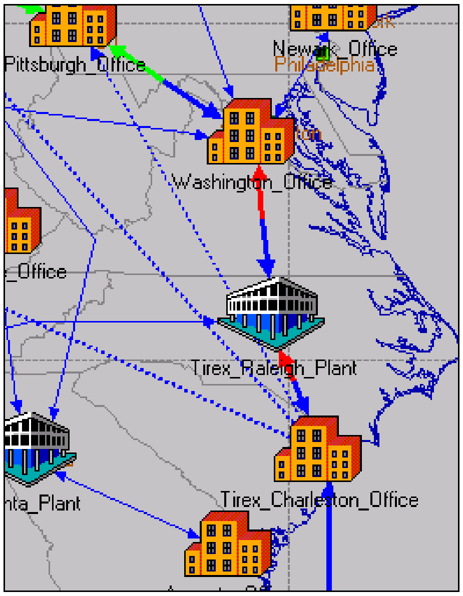

Figure 3-42 Link Color-Coding

7.3 Look for over-utilized links (more than 75% utilized, color-coded red). In this case, there two such links: the link **Tirex_Charleston_Office** ↔**Tirex_Raleigh_Plant** and the link **Tirex_Raleigh_Plant** ↔ **Washington_Office**. To see the exact utilization, place the mouse tip over the link of interest, as shown below.

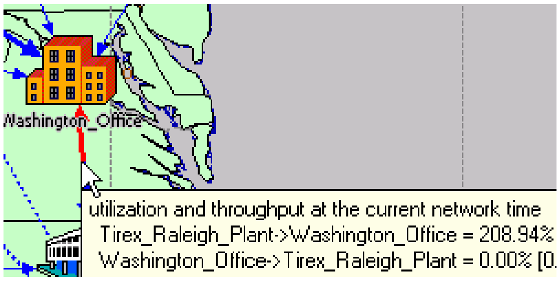

utilization and throughput at the current network time
Tirex_Raleigh_Plant->Washington_Office = 208.94%
Washington_Office->Tirex_Raleigh_Plant = 0.00% [0.

Figure 3-43 Utilization Tip

By placing the mouse tip over each of the links, notice that the links **Tirex_Charleston_Office** ↔**Tirex_Raleigh_Plant** and **Tirex_Raleigh_Plant** ↔ **Washington_Office** are more than 100 % utilized. The link **Tirex_Raleigh_Plant** ↔ **Washington_Office** is intended as a backup link between the two merged networks, the primary link being **Corporate**↔**Tirex_Corporate**.

7.4 We suspect that the traffic is routed through the backup link, leading to its over-loading. To check this hypothesis, select nodes **Tirex_Charleston_Office** and **Detroit_Plant** (press the **Shift** key while selecting both nodes); then choose **Traffic > Visualize > Open Flows Browser**.

7.5 Make sure the **Tirex_Raleigh_Plant**→**Detroit_Plant** is checked; then click **Show Selected**.

7.6 Look at the network to see that the traffic flow between node **Tirex_Raleigh_Plant** and **Detroit_Plant** is routed via the backup link **Tirex_Raleigh_Plant** ↔ **Washington_Office**.

End of Procedure 3-8

Conclusion

Traffic needs to be re-engineered to take the primary link between the two networks, rather than over-load the backup link. The following options are available:

- Reconfigure redistribution and modify the metrics to make the backup link more unlikely to be chosen for routing, and
- Migrate the Tirex network from RIP to EIGRP, such that the merged network will become an all-EIGRP network.

Migrate Tirex Network from RIP to EIGRP

Our example company decided to migrate the Tirex network from running RIP as routing protocol to running EIGRP. In order not to disrupt the traffic over the network, this solution will be implemented in stages.

First, a separate EIGRP process is added to all the routers in the Tirex network. Then, redistribution between the new and old EIGRP processes is configured on border routers.Finally, the new EIGRP process is enabled and the old RIP process is disabled.

Procedure 3-9 Migrating a Network from RIP to EIGRP

1 Duplicate the current scenario.

1.1 Select **Scenario > Duplicate Scenario**.

1.2 Name the new scenario “Merged_Network_Redist_EIGRP”.

2 Enable EIGRP along with RIP on all links in the Tirex network.

2.1 Select the link between **Tirex_Corporate** and **Tirex_New_Orleans** office by clicking on it.

2.2 Right-click on the selected link, and choose **Select Similar Links**.

2.3 Select **Protocols > IP > Routing >Configure Routing Protocols...**.

2.4 Place a checkmark next to EIGRP – leave RIP checked; make sure the Interfaces Across Selected Links option is selected; click **OK**.

2.5 A routing domain visualization appears on the network; notice in the Tirex network, two protocols are running simultaneously: RIP and EIGRP (links will be annotated with the letters ‘E’ from EIGRP and ‘R’ from RIP).

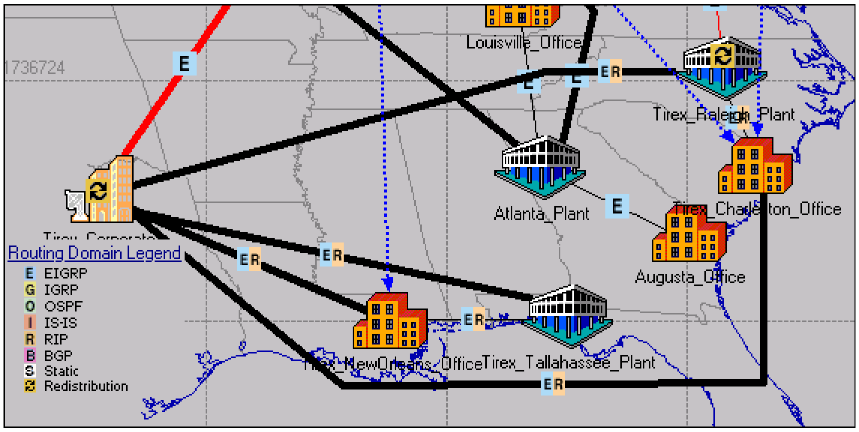

Figure 3-44 Network Showing EIGRP and RIP

2.6 Press **Ctrl+Shift+C** to clear the routing domain visualization.

3 Add a second EIGRP process to all routers in the Tirex network.

3.1 While the links in the Tirex network are still selected, choose **Protocols > EIGRP > Configure Autonomous System...**.

3.2 Enable the **Explicitly Set AS to** radio button and enter **2** as AS number; enable the **Interfaces across selected link(s)** radio button; click **OK**.

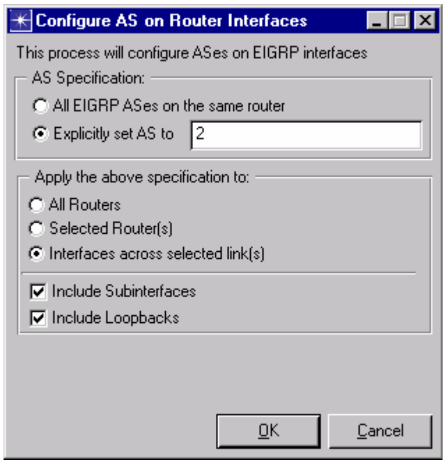

Configure AS on Router Interfaces

This process will configure ASes on EIGRP interfaces

AS Specification:

○ All EIGRP ASes on the same router
● Explicitly set AS to 2

Apply the above specification to:

○ All Routers
○ Selected Router(s)
● Interfaces across selected link(s)

☑ Include Subinterfaces
☑ Include Loopbacks

Figure 3-45 Configure AS on Router Interfaces dialog box

4 Check that the administrative weight of RIP is set to a lower value inside the Tirex network.

4.1 Normally, the routes found by EIGRP take precedence over routes found by RIP in the process of insertion into the IP common route table. The relative precedence is controlled by the administrative weight given to each protocol. When migrating from RIP to EIGRP in the Tirex network, in the first phase, we want to deploy EIGRP without promoting its routes into the common routing table. For that reason, we need to lower the administrative weight of RIP, so the routes found by RIP will always take precedence over the routes found by EIGRP. The routers in the Tirex network have already been configured with a lower administrative weight for RIP than for EIGRP.

4.2 Right-click on the **Tirex_Charleston_Office** and choose **Edit Attributes**; expand the IP attribute category and choose **IP Routing Parameters** then **Administrative Weights**.

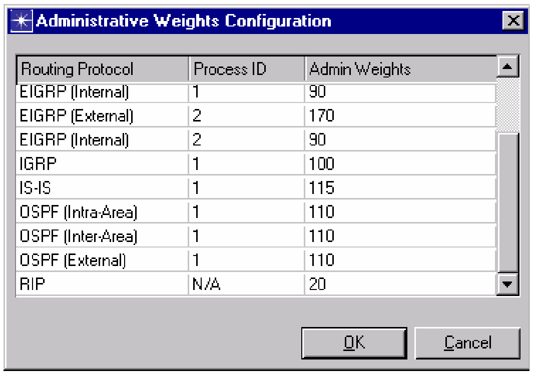

Administrative Weights Configuration

Routing Protocol	Process ID	Admin Weights
EIGRP (Internal)	1	90
EIGRP (External)	2	170
EIGRP (Internal)	2	90
IGRP	1	100
IS-IS	1	115
OSPF (Intra-Area)	1	110
OSPF (Inter-Area)	1	110
OSPF (External)	1	110
RIP	N/A	20

Figure 3-46 Administrative Weights Configuration dialog box

4.3 Notice RIP has a weight of 20, making its routes preferred over routes found by EIGRP, which has an administrative weight of 90.

4.4 Click **Cancel** repeatedly to dismiss the Attribute windows.

5 Check that link utilizations did not change.

5.1 Click the **Configure/Run Simulation** toolbar button.

5.2 Click **Run** to start the simulation.

5.3 Close the Simulation Sequence dialog box after the simulation runs.

5.4 The links should be colored by utilization as before. If they are not, choose **View > Visualize Link Loads > Color by Link Load**.

5.5 Notice the backup link between **Washington_Office** and **Tirex_Raleigh_Plant** is still overloaded.

5.6 To clear the link colors, select **View > Visualize Link Loads > Clear Visualization**.

6 Configure redistribution between the two EIGRP processes.

6.1 Click on the **Tirex_Corporate** node; then right-click on the selected node and select **Select Similar Nodes**. Two nodes will be selected: **Tirex_Corporate** and **Tirex_Raleigh_Plant**.

6.2 Right-click on **Tirex_Corporate** and choose **Edit Attributes**.

6.3 Check the Apply Changes to Selected Objects checkbox in the lower left corner.

6.4 Expand IP Routing Protocols by clicking on the (+) button.

6.5 Choose **EIGRP Parameters** then **Process Parameters**, and double-click on the Process Parameters field in the row corresponding to AS 1.

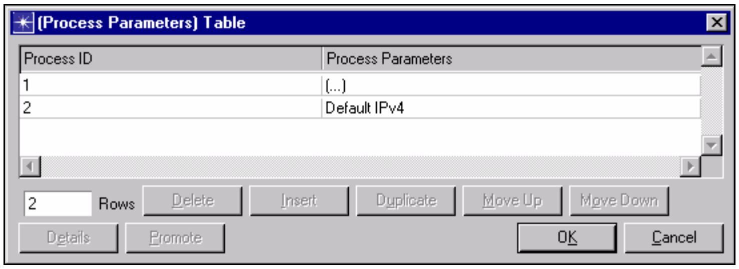

Process ID	Process Parameters
1	(...)
2	Default IPv4

2 Rows

Figure 3-47 Process Parameters Table dialog box

6.6 Choose **Address Family Parameters > Redistribution > Routing Protocols > EIGRP**.

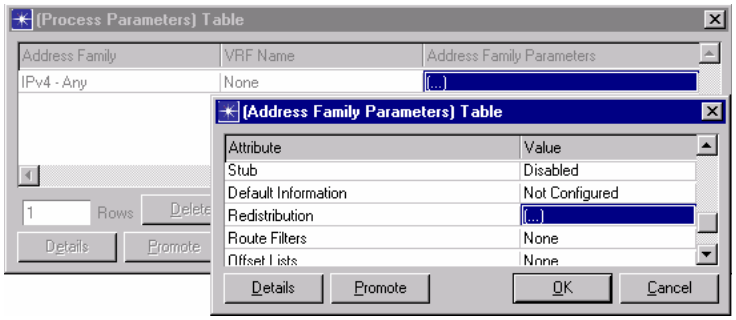

(Process Parameters) Table

Address Family	VRF Name	Address Family Parameters
IPv4 - Any	None	(...)

Rows: 1

(Address Family Parameters) Table

Attribute	Value
Stub	Disabled
Default Information	Not Configured
Redistribution	(...)
Route Filters	None
Offset Lists	None

Figure 3-48 Address Family Parameters Table dialog box

6.7 Make sure that all the settings are as below; by choosing AS Number 2, route redistribution will be enabled from the EIGRP process of AS 2 to the EIGRP process of AS 1.

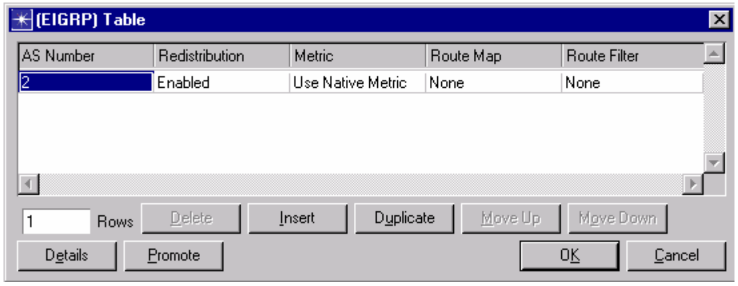

(EIGRP) Table

AS Number	Redistribution	Metric	Route Map	Route Filter
2	Enabled	Use Native Metric	None	None

Rows: 1

Figure 3-49 EIGRP Table dialog box

6.8 Click **OK** to close each open dialog box until you are back at the Process Parameters Table dialog box.

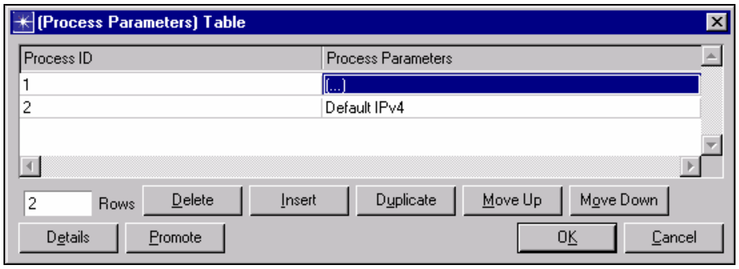

(Process Parameters) Table

Process ID	Process Parameters
1	(...)
2	Default IPv4

Rows: 2

Figure 3-50 Process Parameters Table dialog box

6.9 Double click on the Process Parameters field in the row corresponding to AS 2.

6.10 Choose **Address Family Parameters > Redistribution > Routing Protocols > EIGRP**.

6.11 Make sure all the settings are as below; by choosing AS Number 1, route redistribution will be enabled from the EIGRP process of AS 1 to the EIGRP process of AS 2.

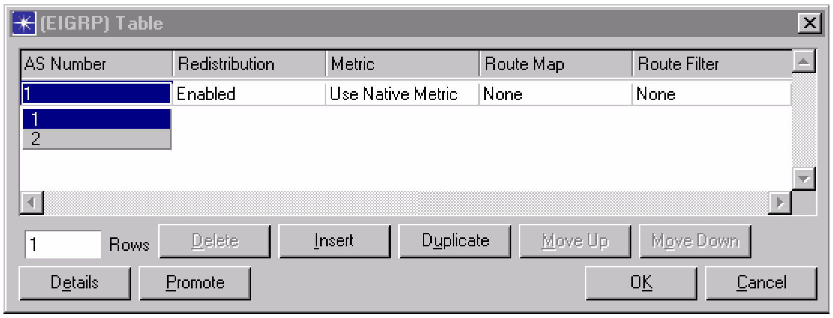

(EIGRP) Table

AS Number	Redistribution	Metric	Route Map	Route Filter
1	Enabled	Use Native Metric	None	None

Figure 3-51 EIGRP Table dialog box

6.12 Click **OK** to dismiss all open attribute windows – a warning will appear to notify you that changes to multiple objects cannot be undone. In this case, you are simultaneously making changes to the two selected routers, in order to enable redistribution between the EIGRP processes running in the two networks. Click **Yes**.

7 Disable RIP processes in the Tirex network.

So far, we have added a second EIGRP process to all nodes in the Tirex network. However, we set the administrative weight of RIP in such a way that any routes found by the newly installed EIGRP processes will be ignored in favor of the old RIP routes. We adopted this approach in order to make a smooth transition from RIP to EIGRP in the Tirex network. While the old RIP routes were still being selected for insertion into the IP forwarding table, the Tirex Company's engineers had had time to debug any issues related to the new EIGRP routes, before these routes came into general use. We know now that EIGRP works fine, and we are ready to replace the RIP routes with EIGRP routes. To achieve this replacement, we will set RIP to an administrative weight higher than that of EIGRP.

7.1 Select all 5 nodes in the Tirex network; to do this, keep the **Ctrl** key pressed while clicking on each of the nodes: **Tirex_Corporate**, **Tirex_New_Orleans_Office**, **Tirex_Tallahasee_Plant**, **Tirex_Charleston_Office**, and **Tirex_Raleigh_Plant.**

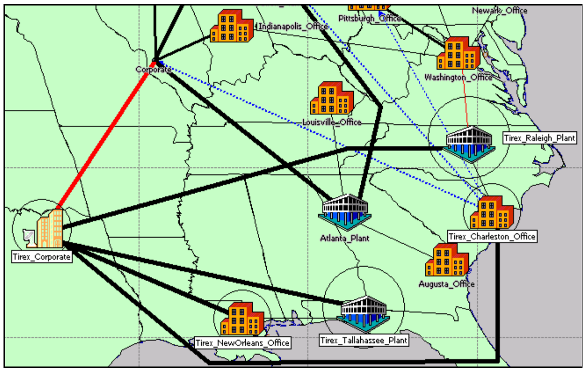

Figure 3-52 Select Tirex Network Nodes

7.2 Right-click on the **Tirex_Corporate** node and select **Edit Attributes**.

7.3 Check the **Apply Changes to Selected Objects** box in the lower left corner.

7.4 Click on the (+) button next to IP and choose **IP Routing Parameters** then **Administrative Weights**.

7.5 Set 500 in the Admin Weights column for RIP.

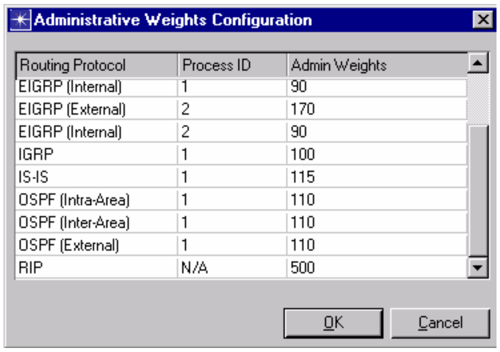

Routing Protocol	Process ID	Admin Weights
EIGRP (Internal)	1	90
EIGRP (External)	2	170
EIGRP (Internal)	2	90
IGRP	1	100
IS-IS	1	115
OSPF (Intra-Area)	1	110
OSPF (Inter-Area)	1	110
OSPF (External)	1	110
RIP	N/A	500

Figure 3-53 Administrative Weights Configuration dialog box

7.6 Click **OK** to dismiss all windows; click **Yes** if you get a warning about changes to multiple objects.

8 Run the simulation.

8.1 Click the **Configure/Run Simulation** toolbar button.

8.2 Click **Run** to start the simulation.

8.3 Close the Simulation Sequence dialog box after the simulation runs.

8.4 Choose **View > Visualize Link Loads > Color by Link Load**; click **OK** to dismiss the description of the link-coloring feature.

End of Procedure 3-9

Results Analysis

Notice the backup link between **Washington_Office** and **Tirex_Raleigh_Plant** is no longer overloaded.

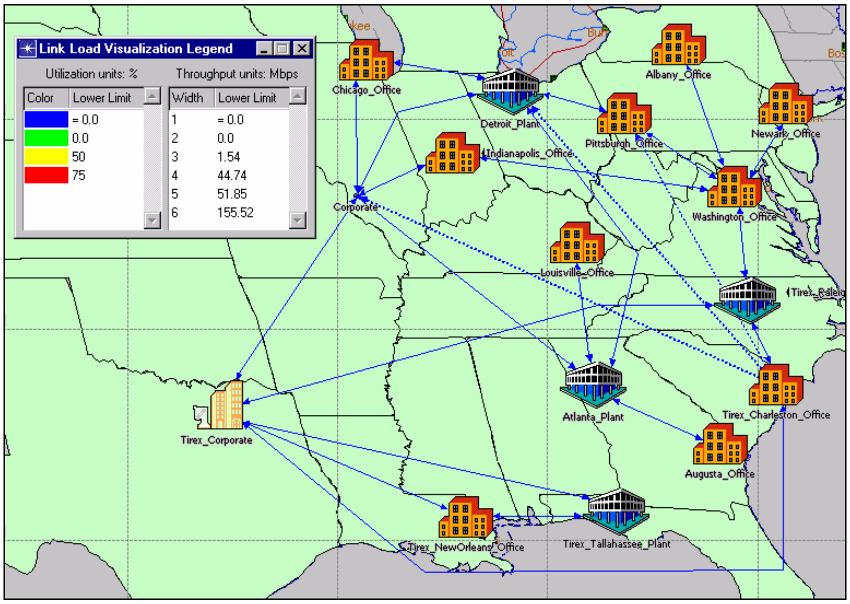

Utilization units: %

Color	Lower Limit
🟦	= 0.0
🟩	0.0
🟨	50
🟥	75

Throughput units: Mbps

Width	Lower Limit
1	= 0.0
2	0.0
3	1.54
4	44.74
5	51.85
6	155.52

Figure 3-54 Washington_Office and Tirex_Raleigh_Plant Backup Link

Conclusion

In the previous example we followed a case of merging two networks running different distance vector protocols: EIGRP and RIP. During the merger, we made use of multiple processes to enable a smooth transition from RIP to EIGRP in one of the networks.

4 Planning and Analyzing IP Multicast Networks

Introduction

IP Multicast is a bandwidth conserving technology that helps reduce traffic by delivering a single stream of information to multiple receivers.

In this chapter, we will:

- Import a simple network with multicast configurations,
- Validate the configurations using NetDoctor,
- Use Virtual Command Line Interface to modify attributes,
- Configure multicast demands,
- Run simulations and visualize the multicast trees and groups,
- Add support for explicit multicast applications on the end nodes, and
- Analyze the impact of the new traffic on the network and visualize link loads and throughputs.

Note—The following examples were presented at OPNETWORK 2004 in Session 1316, Planning and Analyzing IP Multicast Networks. If you do not have access to the files that these procedures use, you can still follow along using the sample screens provided in this user's guide.

There is a set of models with all of the procedural steps completed. You may use these for reference, if you do not wish to do the exercises yourself. These files end in the suffix "_ref" (reference).

Import and Analyze a Multicast Network

A set of router configuration files is provided for an enterprise-sized network running IP multicast. In this section, we will import the configuration files to create a network model and verify its correctness using NetDoctor and analyze multicast routing.

Procedure 4-1 Creating a Network Model from Device Configuration Files

1. Launch NETWARS, if not already opened.
2. From the System Editor's **File** menu, choose **Open Editor**.

3 From the Open Editor drop-down menu, choose **Scenario Builder**, and then click **OK**.

4 Select **File > New Project**. The New Project dialog box displays.

5 In the New Project dialog box, name the new project file `Session_1316_Lab1`, and the phase name `Background_Traffic`, and then click **OK**.

6 Choose **Topology > Import > Device Configuration Files...**.

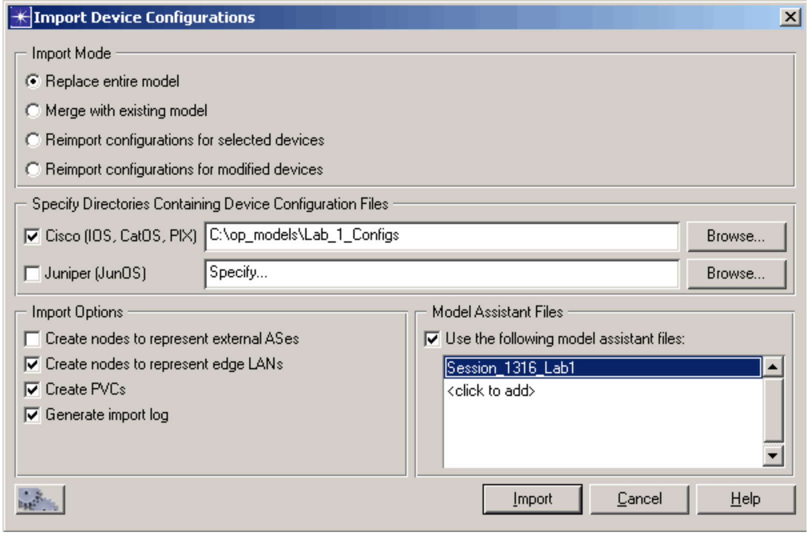

Import Device Configurations

Import Mode

● Replace entire model
○ Merge with existing model
○ Reimport configurations for selected devices
○ Reimport configurations for modified devices

Specify Directories Containing Device Configuration Files

☑ Cisco (IOS, CatOS, PIX) C:\op_models\Lab_1_Configs
☐ Juniper (JunOS) Specify...

Import Options

☐ Create nodes to represent external ASes
☑ Create nodes to represent edge LANs
☑ Create PVCs
☑ Generate import log

Model Assistant Files

☑ Use the following model assistant files:
Session_1316_Lab1
<click to add>

Figure 4-1 Import Device Configurations dialog box

7 Specify folders for device configuration files:

7.1 Select the checkbox for **Cisco (IOS, CatOS, PIX)**.

7.2 Click **Browse** for Cisco (IOS, CatOS, PIX) and select the following folder:

```
C:\op_models\Lab_1_Configs
```

7.3 De-select the checkbox for **Juniper (JunOS)**.

8 Specify import options:

8.1 Select the checkbox for **Create nodes to represent edge LANs**. This option creates an edge LAN node for each unconnected, active interfaces. Edge LANs can be used as a traffic source or destination.

8.2 Select the checkbox for **Create PVCs**. This option creates a full mesh of PVCs for frame relay interface that belong to the same subnet.

8.3 Select the checkbox for **Generate import log**.

9 Click *<Click to add>* and choose *Session_1316_Lab1* (this Model Assistant file is used to specify connectivity and location information).

10 Click **Import**.

11 **Import Summary (Concise)** appears; observe that there are no skipped configuration files. Click **Close**.

12 Save the project:

12.1 Choose **File > Save**.

12.2 Verify that file name is `Session_1316_Lab_1`, and click **Save**.

13 Hide all the PVCs (demands) to clear up the network topology. The imported network should resemble the following figure.

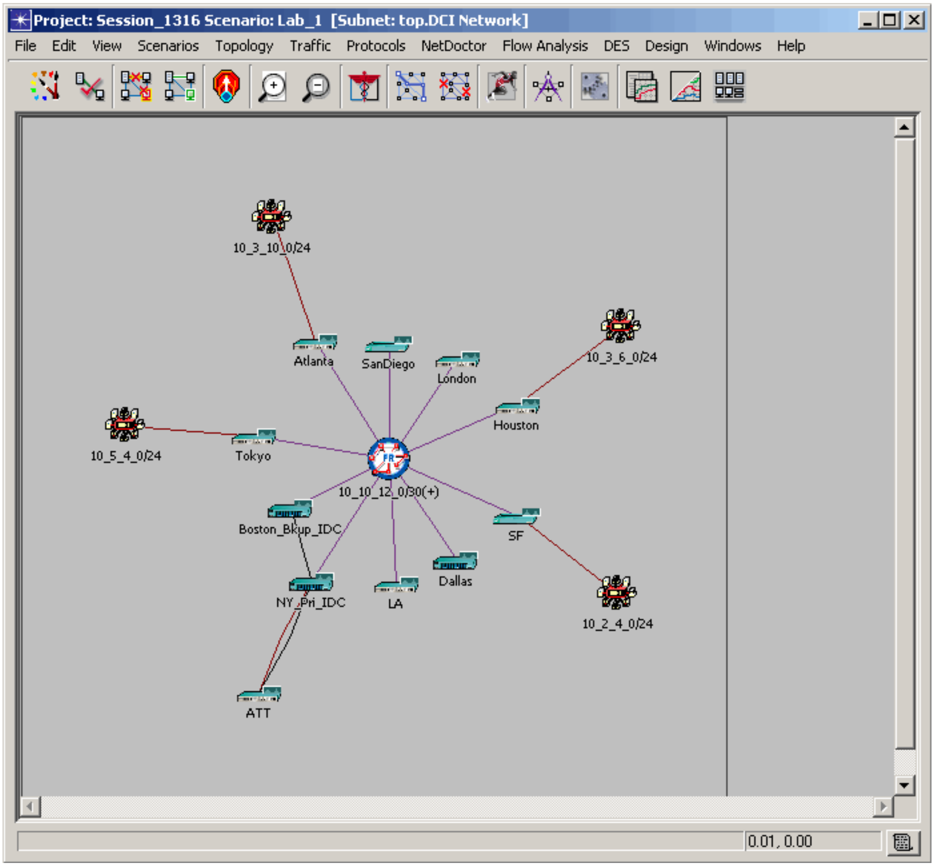
Project: Session_1316 Scenario: Lab_1 [Subnet: top.DCI Network]
File Edit View Scenarios Topology Traffic Protocols NetDoctor Flow Analysis DES Design Windows Help
10_3_10_0/24
Atlanta
SanDiego
London
10_3_6_0/24
Houston
10_5_4_0/24
Tokyo
10_10_12_0/30(+)
Boston_Bkup_IDC
SF
Dallas
NY_Pri_IDC
LA
10_2_4_0/24
ATT
0.01, 0.00

Figure 4-2 Imported Network

The import process is now complete.

End of Procedure 4-1

Understand the Imported Network

In this section, we explore and analyze the imported network model using Flow Analysis and user-defined reports.

Procedure 4-2 Exploring the Network Model

1 Select **Scenarios > Summary Tables > User-Defined Reports > Generate Report from Template…**. The Generate User-Defined Report dialog box displays.

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 one table.

2 Expand **Multicast** and select **Multicast Groups** and **RP Configuration** items listed as reports to be generated.

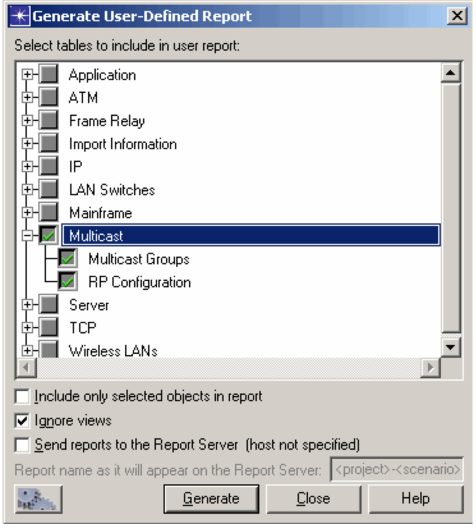

Generate User-Defined Report

Select tables to include in user report:

- □ Application
- □ ATM
- □ Frame Relay
- □ Import Information
- □ IP
- □ LAN Switches
- □ Mainframe
- ☑ Multicast
  - ☑ Multicast Groups
  - ☑ RP Configuration
- □ Server
- □ TCP
- □ Wireless LANs

□ Include only selected objects in report

☑ Ignore views

□ Send reports to the Report Server (host not specified)

Report name as it will appear on the Report Server: <project>-<scenario>

Figure 4-3 Generate User-Defined Report dialog box

3 Click **Generate**. The View Results dialog box displays.

4 Browse the View Results dialog box:

4.1 Expand **Multicast** and select “Multicast Groups”. Observe that there is one multicast group with address **236.1.1.2**.

4.2 Click **Show**. Members of this group are routers **Atlanta, Boston, Houston, SF and Tokyo.**

4.3 Select **RP Configuration** and observe that Atlanta has Auto-RP status enabled for the group identified by the ACL 31.

4.4 For either report, you can view the full table by clicking **Show** or double-clicking on the selected item on the tree-view.

4.5 Close all the tables and the View Results dialog box

End of Procedure 4-2

Enable Multicast on End-Nodes

The edge LANs created are not multicast-enabled. If we need them to send or receive multicast traffic, we need to enable them explicitly.

Procedure 4-3 Configuring Multicast Traffic Flows

1 All the end-nodes (LANs) should be multicast-enabled:

1.1 Select the LAN node connected to *Tokyo* (*10_5_4_0/24*).

1.2 Right-click on it and choose **Select Similar Nodes**.

1.3 Select **Protocols** > **IP > Multicast > Enable Multicasting on Selected Hosts**.

This option essentially enables the attribute **IP/ IP Host Parameters / Multicast Mode**.

1.4 Make sure you get a confirmation saying "*Multicasting has been enabled on 4 hosts"*.

End of Procedure 4-3

Add Traffic to the Imported Network

Procedure 4-4 Adding Traffic to the Network

1 Click on the **Object Palette** tool bar button.

2 From the pull-down menu, choose the "**demands**" palette.

3 Select the *ip_mcast_traffic_flow* from the palette.

4 After selecting the above demand:

4.1 Click on the LAN node connected to **SF** (*10_2_4_0/24)* to initiate the demand-creation operation.

4.2 Right-click on the empty project area above the node, and choose **Finish Demand Definition** to create a demand.

4.3 Right-click again on the empty project area and choose **Abort Demand Definition** to finish the demand creation operation.

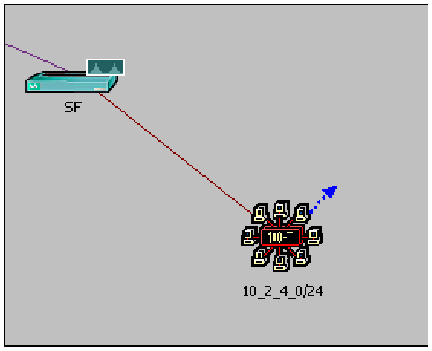

Figure 4-4 Creating a Demand

5 Close the Object Palette.

6 Specify the destination multicast address for the demand:

6.1 Right-click on the created demand and choose **Edit Attributes**.

6.2 For the Destination Address, enter **236.1.1.2**. This is the multicast address to which the routers have been configured to listen.

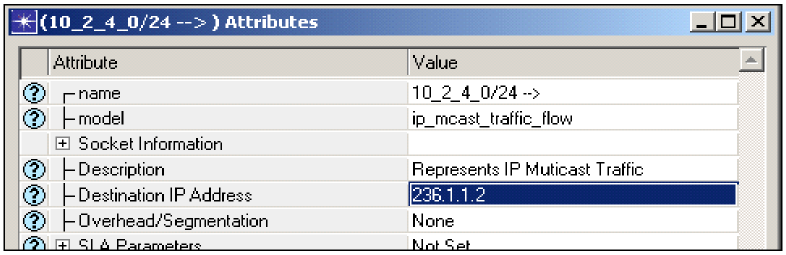
(10_2_4_0/24 -->) Attributes

Attribute	Value
name	10_2_4_0/24 -->
model	ip_mcast_traffic_flow
⊞ Socket Information	
Description	Represents IP Muticast Traffic
Destination IP Address	236.1.1.2
Overhead/Segmentation	None
⊞ SLA Parameters	Not Set

Figure 4-5 Specifying the Destination Address

7 Specify the traffic characteristics:

7.1 Double-click on the **Traffic (bits/second)** attribute value. The **Traffic (bits/second) Attribute Profile** dialog box appears.

7.2 Click on the **bits/second** column on the first row and enter 100000.

7.3 Click on the **seconds** column on the second row and enter 150.

7.4 Enter 300000 on the second row of the **bits/second** column.

7.5 Finally in the third row enter **3600** in the seconds column and **300000** in the **bits/second** column.

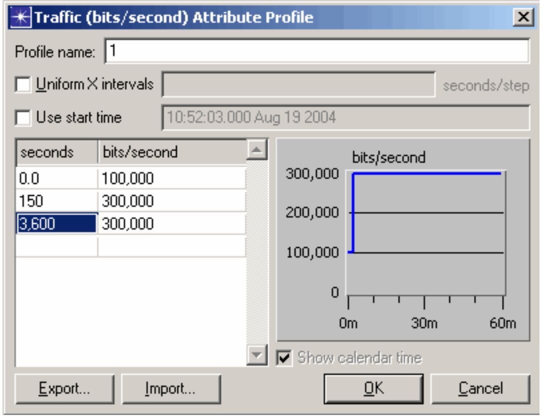

Traffic (bits/second) Attribute Profile

Profile name: 1

□ Uniform X intervals seconds/step

□ Use start time 10:52:03.000 Aug 19 2004

seconds	bits/second
0.0	100,000
150	300,000
3,600	300,000

☑ Show calendar time

Figure 4-6 Traffic (bits/second) Attribute Profile dialog box

7.6 Click **OK** to accept the changes.

Here, we are configuring a step-sized stream of background traffic at the rate of nearly 300 Kbps for duration of 3450 seconds (~1 hour).

We have intentionally decreased the traffic during the first 150 seconds because we want to capture the behavior during the steady state that occurs after the switching from shared tree to shortest-path tree. Thus the switching does not affect the average and peak results for link throughput.

This transitional state occurs because, until a few first packets are sent from the source to the rendezvous point, the rendezvous point (as well as the other destinations) will receive duplicate messages on SPT and ST. This would increase the link throughput for a short duration.

7.7 Double-click on **Traffic (packets/second)**. The **Traffic (packets/second) Attribute Profile** dialog box appears.

7.8 Click on **packets/second** column on the first row and enter 100.

7.9 Click on the **seconds** column on the second row and enter 150.

7.10 Enter 250 on the second row of **packets/second**.

7.11 Enter 3600 and 250 in the third rows for **seconds and packets/second** respectively.

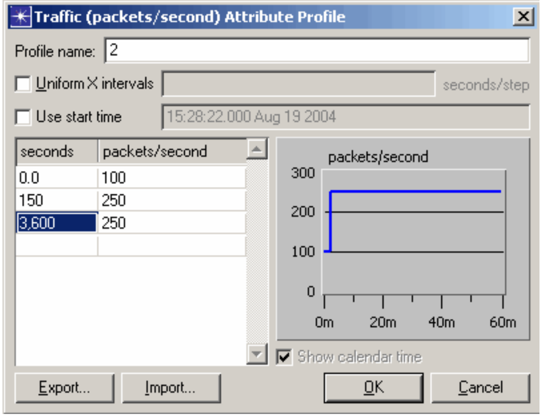

Traffic (packets/second) Attribute Profile

Profile name: 2

□ Uniform X intervals seconds/step

□ Use start time 15:28:22.000 Aug 19 2004

seconds	packets/second
0.0	100
150	250
3,600	250

☑ Show calendar time

Figure 4-7 Traffic (packets/second) Attribute Profile dialog box

7.12 Click **OK** twice to save all changes.

We are specifying the number of packets that must be generated per second for the traffic we configured (300 Kbps/second).

Note that in the first attribute we specified the traffic to be generated in bits/sec and in the second attribute we specify the packets to be generated for the configured amount of traffic.

End of Procedure 4-4

Validate the Network Model Using NetDoctor

Procedure 4-5 Detecting Configuration Errors

1 Choose **NetDoctor > Configure/Run NetDoctor**.

2 Select all the rules in the **Rule Suites / IP Multicast**.

3 Click **Run**.

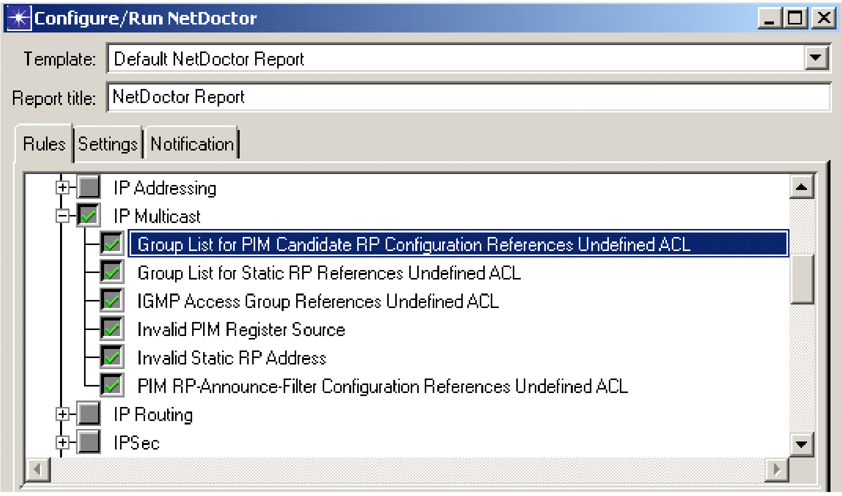
Configure/Run NetDoctor
Template: Default NetDoctor Report
Report title: NetDoctor Report
Rules
Settings
Notification
IP Addressing
IP Multicast
Group List for PIM Candidate RP Configuration References Undefined ACL
Group List for Static RP References Undefined ACL
IGMP Access Group References Undefined ACL
Invalid PIM Register Source
Invalid Static RP Address
PIM RP-Announce-Filter Configuration References Undefined ACL
IP Routing
IPSec

Figure 4-8 Configure/Run NetDoctor dialog box

NetDoctor shows one configuration warning for the rule **Group List for PIM Candidate RP Configuration References Undefined ACL.**

4 Click on the warning on the left frame and observe the detailed message on the right frame.

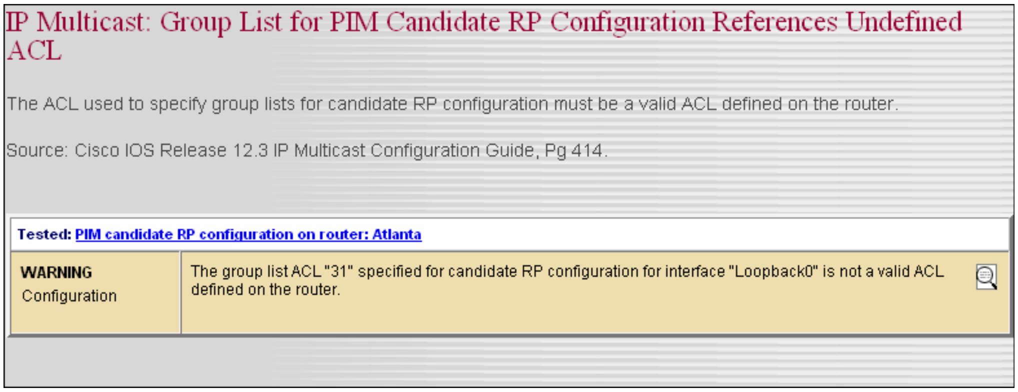

IP Multicast: Group List for PIM Candidate RP Configuration References Undefined ACL

The ACL used to specify group lists for candidate RP configuration must be a valid ACL defined on the router.

Source: Cisco IOS Release 12.3 IP Multicast Configuration Guide, Pg 414.

Tested: PIM candidate RP configuration on router: Atlanta

WARNING Configuration	The group list ACL "31" specified for candidate RP configuration for interface "Loopback0" is not a valid ACL defined on the router.

Figure 4-9 NetDoctor Warning

This could be the cause of un-routable demands in Flow Analysis. First, let us verify the problem by looking into the concerned attribute.

5 Right-click on the node **Atlanta** and choose **Edit Attributes**. Observe the following attribute under the group **IP Multicasting**: **PIM Parameters > Auto-RP Configuration > Candidate RP Configuration > Group Filter Configuration > Group Filter ACL**. Note that the value is **31**.

6 Cancel out all the way to return to the top-level attributes.

7 Now observe the standard ACL that the above attribute refers to: **IP. IP Routing Parameters. Standard ACL Configuration**. Note that we see two ACLs, **10** and **30.**

8 Look in the **List Configuration** attribute for ACL 30. Note that this refers to the multicast address 236.1.1.2 that our routers wanted to join.

MISCONFIGURATION: We see that instead of referring to ACL 30, we have mistakenly referred to ACL 31 when defining the Candidate RPs.

The reason why there were no routes is because we had configured Atlanta to be an RP for the group that is defined by ACL 31. When flow analysis found no such ACL, it ignored the setting. This led to a multicast group (236.1.1.2) without any RP configuration.

There are three ways to fix this error:

8.1 Modify the specified attribute directly.

8.2 Modify the configuration file and do a reimport.

8.3 Modify the specified attribute using the Virtual Command Line Interface (or Virtual CLI.)

In the following procedure, we will use Virtual CLI to fix this misconfiguration.

9 Click **Cancel** to close all the attribute dialog boxes.

End of Procedure 4-5

Use Virtual CLI to Fix Configuration Errors

The Virtual Command Line Interface (or Virtual CLI) emulates Cisco's CLI so that Cisco configuration commands can be entered for OPNET models. This interface is available only for OPNET node models created from Cisco IOS and CatOS configurations.

Procedure 4-6 Using Virtual CLI to Fix Errors

1 The command that you enter on a router to change the above attribute is:

```
ip pim send-rp-announce interface-type interface-number scope
ttl-value [group-list access-list]
```

Based on the Standard ACLs configured on the node and the multicast configuration, the correct command to be entered is:

```
ip pim send-rp-announce Loopback 0 scope 10 group-list 30
```

2 Right-click on the router **Atlanta** and select **Open Virtual CLI…**

The dialog box that appears is the virtual command line interface; you can enter Cisco commands as you would enter them on the real Cisco device.

3 At any point you can:

3.1 Make use of the auto-fill or auto-complete feature by pressing the **Tab** key.

3.2 Display the list of supported commands by typing a ***?*** (question mark).

4 Press **Enter** and type ***en*** to enter the enable mode.

5 Type ***config t*** ↵ to enter configuration commands from the terminal.

6 Type in the command that we identified to fix the error.

```
ip pim send-rp-announce Loopback 0 scope 10 group-list 30 ↵
```

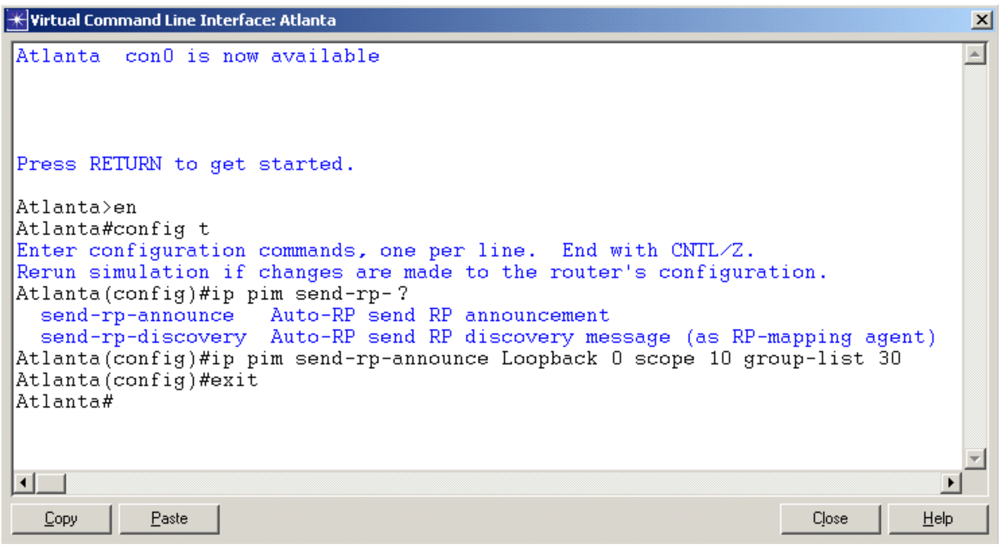

```
Atlanta  con0 is now available



Press RETURN to get started.

Atlanta>en
Atlanta#config t
Enter configuration commands, one per line.  End with CNTL/Z.
Rerun simulation if changes are made to the router's configuration.
Atlanta(config)#ip pim send-rp-?
  send-rp-announce   Auto-RP send RP announcement
  send-rp-discovery  Auto-RP send RP discovery message (as RP-mapping agent)
Atlanta(config)#ip pim send-rp-announce Loopback 0 scope 10 group-list 30
Atlanta(config)#exit
Atlanta#
```

Figure 4-10 Virtual CLI Window

7 Type ***exit*** to come out of the configuration mode.

8 Click **Close** to come out of the Virtual Command line interface.

9 Verify that the attribute was indeed changed to 30.

IP. PIM Parameters. Auto-RP Configuration. Candidate RP Configuration. Group Filter Configuration. Group Filter ACL

10 Rerun NetDoctor to verify that the error was cleared:

10.1 Choose **NetDoctor > Configure/Run NetDoctor**.

10.2 Make sure all the rules in the **Rule Suites / IP Multicast** is selected.

10.3 Click **Run**.

Note that NetDoctor reports zero errors and zero warning messages in the web report.

End of Procedure 4-6

Visualize Multicast Demands and Trees

Procedure 4-7 Visualizing Multicast Demands and Trees

1 Rerun the simulation:

1.1 Select **Flow Analysis** > **Configure/Run Flow Analysis…**

1.2 Make sure the **Start time** and the **Stop time** differs at least by an hour.

1.3 In the Protocol Settings, make sure **Consider all unconnected interfaces** is checked.

1.4 Make sure **Display flow analysis log** is *unchecked*.

1.5 Click **Run**.

2 Select **Traffic > Visualize > Open Flows Browser**.

3 Expand the node 10_2_4_0/24 and click on the demand *10_2_4_0/24*→. You will see a list of routes displayed on the right-hand pane in the window.

4 Click on each of the routes to observe the path taken to each of the destination nodes (that is, the group members).

5 Observe that all the routes taken by the demand is the shortest path to the destination, indicating that it has taken the SPT (Shortest Path Tree). By default, all the routers—even those configured to run PIM-SM—will switch to shortest path after the first multicast packet.

We have successfully fixed the un-routable demands problem using NetDoctor and Virtual CLI.

6 View multicast trees:

6.1 Select the LAN node connected to **SF** *10_2_4_0/24,* and choose **Show Multicast Routes From Selected Node…**

6.2 In the View Multicast Route dialog box, make sure that **236.1.1.2** is selected as the Multicast Group Address and that the **Source-based tree** check box is checked.

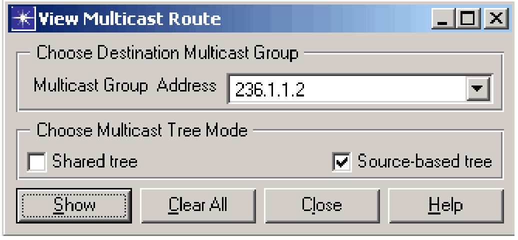

View Multicast Route

Choose Destination Multicast Group

Multicast Group Address 236.1.1.2

Choose Multicast Tree Mode

☐ Shared tree ☑ Source-based tree

Figure 4-11 View Multicast Route dialog box

6.3 Click **Show**.

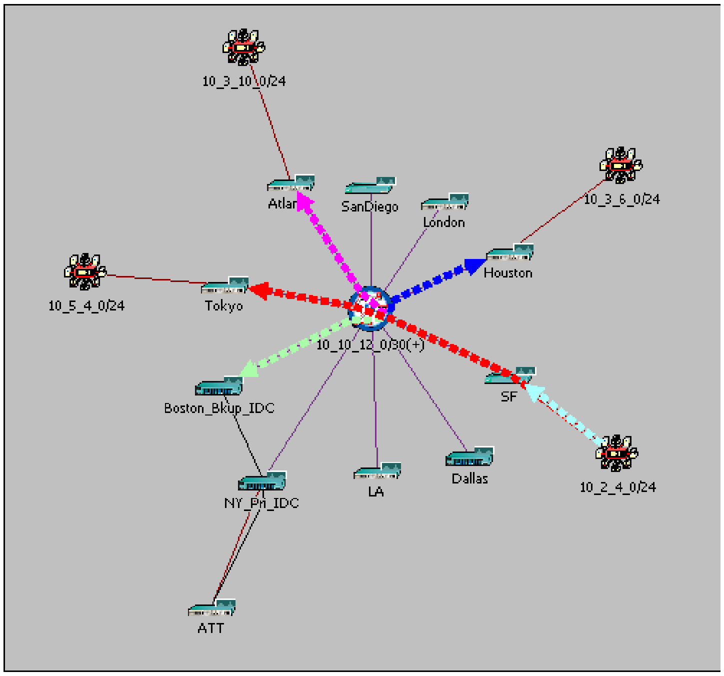

Figure 4-12 Showing Multicast Routes

6.4 Observe that the route from SF takes the shortest path to all the group members i.e. Tokyo, Atlanta, Houston and Boston.

6.5 Click **Clear All**.

6.6 Now check the checkbox for **Shared tree**.

6.7 Click **Show**.

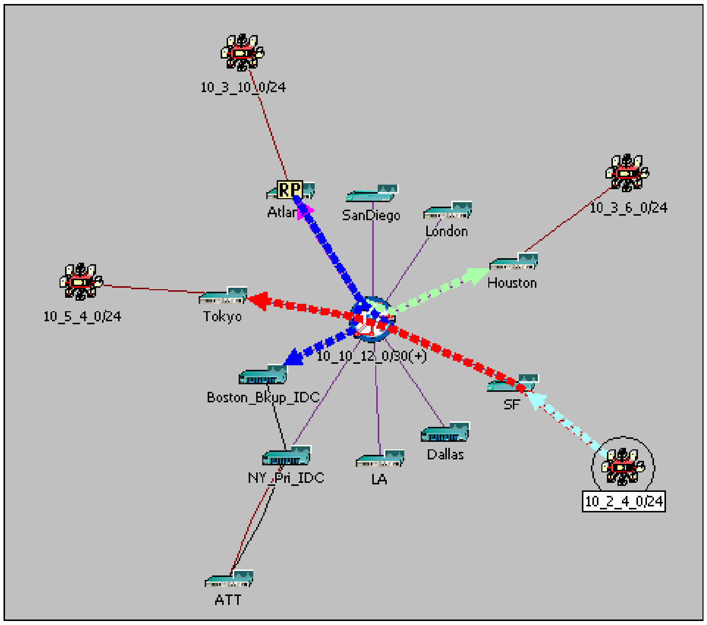

Figure 4-13 Showing Shared Tree Routes

6.8 Observe the small **RP** icon on the node **Atlanta**.

Note that (except for Tokyo) all the traffic goes to Atlanta and then to the respective nodes; we will later investigate why Tokyo traffic goes via the shortest path instead of taking the path through Atlanta.

6.9 To observe all traffic going to Atlanta, click **Close** in the View Multicast Route dialog box (make sure you still see the paths in the network model).

6.10 Click on the paths or the link between the central frame-relay cloud and node **Atlanta**.

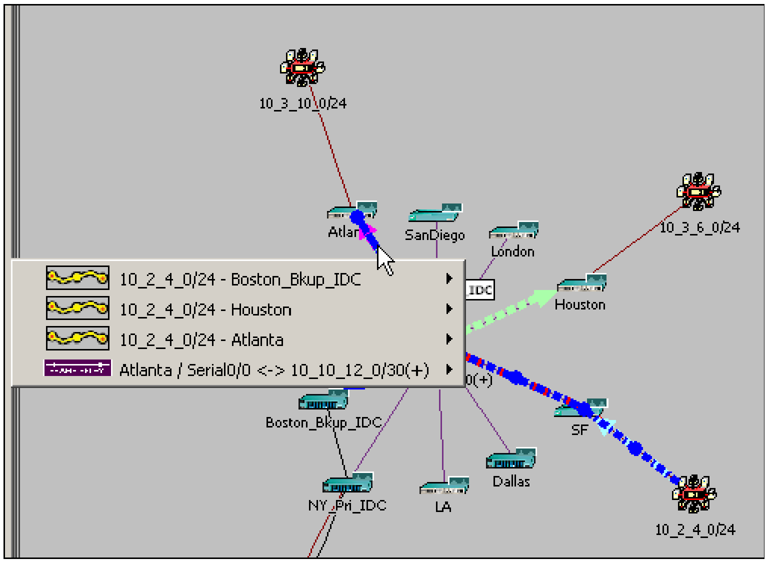

Figure 4-14 Showing Traffic Through Atlanta

6.11 Note that you see three paths from 10_2_4_0/24 destined to nodes Boston and Houston traveling through Atlanta.

We saw the shared tree from **SF** going directly to **Tokyo** because the normal unicast route between **Tokyo** and **Atlanta** (RP) passes through **SF.** When **SF** sends traffic to **Atlanta**, it comes back to **SF;** hence from that point on, **SF** learns to bypass Atlanta and send traffic to Tokyo directly, even if it is a shared tree.

End of Procedure 4-7

Observe Link Usage and Utilization

We will study the link utilization in the network due to the current levels of multicast traffic.

Procedure 4-8 Observing Link Usage and Utilization

1 Select **View** > **Visualize Link Loads > Color by Link Load ...**.

2 In the Color Links by Load dialog box, make sure that **Flow Analysis** is chosen for the first pull-down menu and **peak utilization and throughput for each link** for the second pull-down menu.

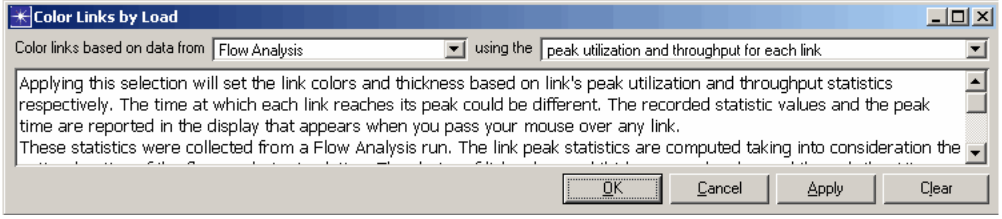

Color links based on data from: Flow Analysis
using the: peak utilization and throughput for each link

Applying this selection will set the link colors and thickness based on link's peak utilization and throughput statistics respectively. The time at which each link reaches its peak could be different. The recorded statistic values and the peak time are reported in the display that appears when you pass your mouse over any link.
These statistics were collected from a Flow Analysis run. The link peak statistics are computed taking into consideration the

Figure 4-15 Color Links by Load dialog box

3 Click **Apply** to see link widths and colors change according to the throughput and peak utilization.

4 Note that the only link that is heavily utilized is the one connecting **SF** and the Frame relay cloud *10_10_12_0/30(+)*.

5 Move your mouse over the links to observe the peak utilization and throughput for that and the other links.

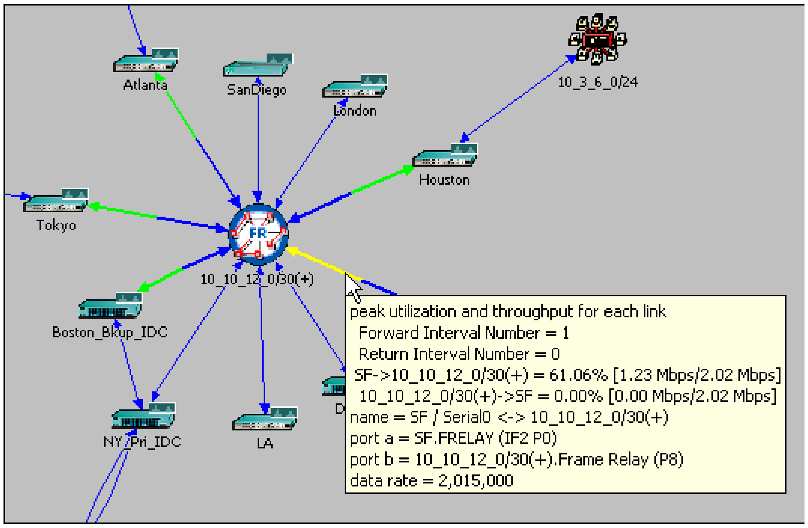

```
peak utilization and throughput for each link
  Forward Interval Number = 1
  Return Interval Number = 0
 SF->10_10_12_0/30(+) = 61.06% [1.23 Mbps/2.02 Mbps]
  10_10_12_0/30(+)->SF = 0.00% [0.00 Mbps/2.02 Mbps]
name = SF / Serial0 <-> 10_10_12_0/30(+)
port a = SF.FRELAY (IF2 P0)
port b = 10_10_12_0/30(+).Frame Relay (P8)
data rate = 2,015,000
```

Figure 4-16 Peak Utilization and Throughput Tip

End of Procedure 4-8

Conclusions

The Device configuration import allows you to import a network topology using device configuration files. Using NetDoctor, you can validate your network and detect any configuration problems. Virtual command line interface enables you to fix the error with a user-interface that emulates Cisco CLI.

You used Flow Analysis to visualize the multicast trees and the paths taken by the demands. You visualized the difference between the shortest path tree and the shared tree (RPT).

In the next section, you will learn how to configure explicit multicast applications like video-conferencing on this imported network.

Adding Explicit Shared Multicast Applications to a Network

This section walks you through all the configurations required to configure the explicit multicast traffic on a network.

In the following procedure, we will send video traffic between **10_2_4_0/24** (source) and **10_3_6_0/24** (destination). In the process, we will learn how to configure explicit multicast traffic on the network using the application objects, and run a discrete event simulation to study the impact of the new traffic on the network.

Procedure 4-9 Adding Explicit Applications to a Network

1 Launch NETWARS, if not already opened.

2 From the System Editor's **File** menu, choose **Open Editor**.

3 From the Open Editor drop-down menu, choose **Scenario Builder**, and then click **OK**.

4 Select **File > Open Project**. The Open Project dialog box displays.

5 In the Open Project dialog box, select the project file `Session_1316_Lab2` in folder `C:\op_models`, and then click **Open**.

5.1 Make sure the current scenario is *Explicit_Applications*.

6 Specify the application:

6.1 Right-click on the **Application Config** node and select **Edit Attributes**.

6.2 Verify that there is a single application called **Video**.

6.3 Observe the specification of the application looking at the attributes under **Description** / **Video Conferencing**.

- Note that the **Symbolic Destination Name** for the application is set as **Multicast Destination**.
- The **Frame Interarrival Time Information** attribute specifies the interarrival time between the frames (e.g 10 Frames /second) using a distribution.
- The **Frame Size Information** attribute specifies the frame size of the incoming and the outgoing video streams.

6.4 Click **Cancel**, all the way to the node level.

6.5 Right-click on the Profile Config node and select **Edit Attributes**.

6.6 Verify that there is a single profile called **Video**.

6.7 Click on the **Applications** attribute in that row to observe more details about this profile.

6.8 Click **Cancel**, all the way to the node level.

7 Configure the application:

Note that in the application definition we have defined a Symbolic Destination called **Multicast Destination.** This symbolic name needs to be resolved or mapped onto a real destination address (in our case, a multicast destination address).

7.1 Right-click on the LAN object attached to **SF** i.e. 10_2_4_0/24, and then choose **Edit Attributes**.

7.2 Expand the group **Applications** and double-click on the attribute **Application: Destination Preferences.**

7.3 Observe that there is a single row with **Video** set as the value for the attribute **Application** and that **Multicast Destination** is set as the **Symbolic Name**.

7.4 Click on the **Actual Name** attribute value and note that **Name** is configured to **236.1.1.2**.

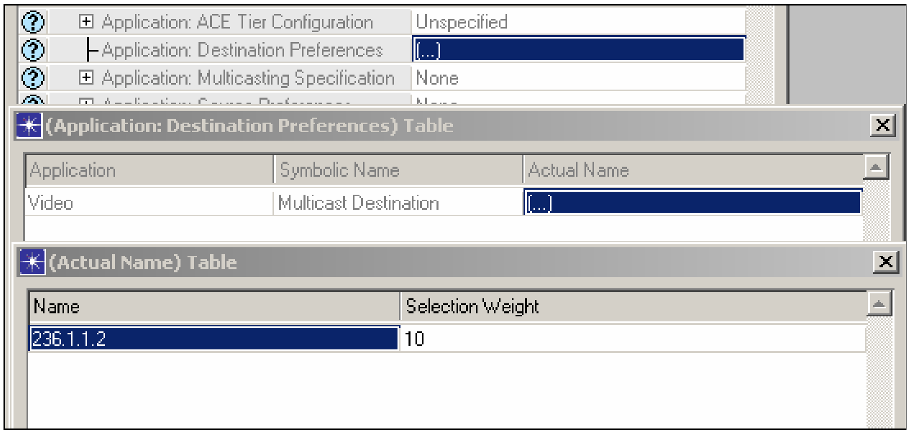

⊞ Application: ACE Tier Configuration	Unspecified
Application: Destination Preferences	(...)
⊞ Application: Multicasting Specification	None
⊞ Application: Source Preferences	None

(Application: Destination Preferences) Table

Application	Symbolic Name	Actual Name
Video	Multicast Destination	(...)

(Actual Name) Table

Name	Selection Weight
236.1.1.2	10

Figure 4-17 Actual Name Table Attribute dialog box

7.5 Click **Cancel** twice to return to the top-level attribute structure.

7.6 Click on the **Application: Supported Profiles** attribute.

7.7 Observe that the **Profile Name** is set as **Video**.

7.8 Click **Cancel** to come back to the top-level attributes.

7.9 Notice that the value for the attribute **Application: Transport Protocol** is set as **UDP**. (This is because OPNET currently supports only UDP-based multicast applications).

8 Join the multicast groups:

Thus far we configured the application and the resolved the destination address. Now we need to specify which end-nodes are going to listen/receive this multicast video traffic.

8.1 Right-click on the LAN object attached to Houston **(**i.e. 10_3_6_0/24) and choose **Edit Attributes**.

8.2 Expand the group **Applications** and double-click on the attribute **Application: Multicasting Specification**.

8.3 Notice that the attribute **Application** is set to **Video**.

8.4 Double-click on the attribute **Membership Addresses** to see **236.1.1.2** set as the **Supported Multicast Addresses**.

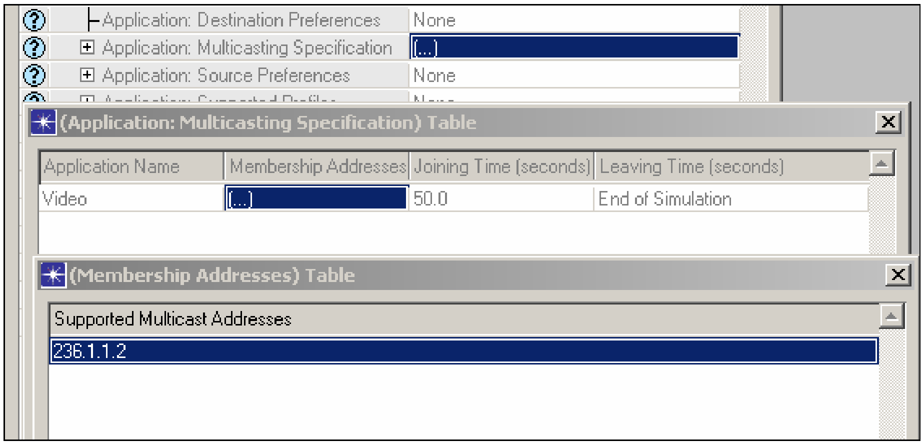

Application: Destination Preferences	None
Application: Multicasting Specification	(...)
Application: Source Preferences	None

(Application: Multicasting Specification) Table

Application Name	Membership Addresses	Joining Time (seconds)	Leaving Time (seconds)
Video	(...)	50.0	End of Simulation

(Membership Addresses) Table

Supported Multicast Addresses
236.1.1.2

Figure 4-18 Membership Addresses Table Attribute dialog box

8.5 Click **Cancel** twice to come back to the top-level attributes.

8.6 Double-click on the attribute value for **Applications: Supported Services.**

8.7 Note that application **Video** is set as one of the supported applications on this node.

8.8 Click **Cancel** to come back to the top-level attribute structure.

8.9 Make sure that the attribute **Application: Transport Protocol,** is set to UDP.

8.10 Click **Cancel** to return to the network.

9 Choose statistics:

9.1 Right-click in the project workspace and select **Choose Individual DES Statistics**.

9.2 Expand the **link statistics** and under **point-to-point** statistics, make sure the statistics **throughput (bits/sec) and utilization** in both directions have been selected.

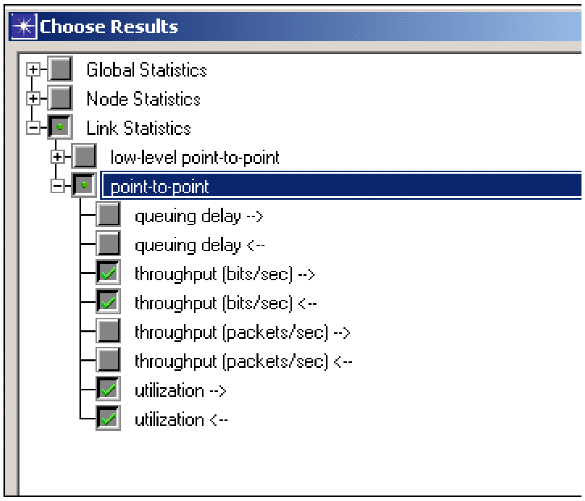

Figure 4-19 Choose Results dialog box

10 Configure and run Discrete Event Simulation:

10.1 Click the **Configure/Run Simulation** button.

Note that the **Duration** is set to “5.0” and that the corresponding unit is set to **minute(s).**

10.2 Click **Run**. The simulation takes about 1-2 minutes to complete.

11 Perform link visualization:

11.1 Close any visible graphs or View Results dialog box.

11.2 Select **View > Visualize Link Loads > Color Link by Link Load ….**

11.3 In the Color Links by Load dialog box, make sure **Discrete Event Simulation** is chosen for the first pull-down menu and **peak utilization and throughput for each link** for the second-pull down menu.

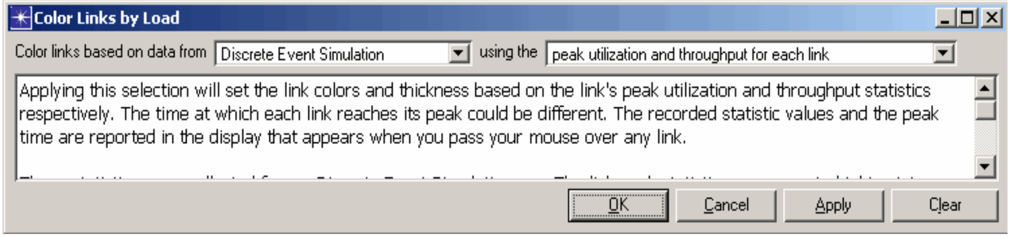

Color links based on data from Discrete Event Simulation using the peak utilization and throughput for each link

Applying this selection will set the link colors and thickness based on the link's peak utilization and throughput statistics respectively. The time at which each link reaches its peak could be different. The recorded statistic values and the peak time are reported in the display that appears when you pass your mouse over any link.

Figure 4-20 Color Links by Load dialog box

11.4 Click **Apply** to see link width and the colors change according the throughput and peak utilization.

11.5 Note that the link between **SF** and the Frame Relay cloud is now shown in red, indicating that the utilization has exceeded the 75 % threshold due to our video conferencing multicast traffic.

You can change the option to **average utilization and throughput for each link** to view the average link health.

End of Procedure 4-9

Conclusions

Explicit multicast applications can be configured on the end-nodes using the application definition object and the profile definition object. Modeling explicit traffic allows us to calculate packet end-to-end delay and thus evaluate the application performance effectively.

Link visualization allows us to visualize the peak and average utilization and throughput for the duration of simulation.