Using the Analyzer

For purposes of this chapter, it is assumed that you have read Chapter 1, “Analyzer Overview,” which describes the concepts and principles essential to understanding and using the Analyzer.

This chapter tells you how to use the Analyzer by means of the graphical user interface (GUI) of the Display module. The Display modules, which can be installed on one or more host workstations or PCs in your network, incorporates an extensive menu system for invoking Analyzer functionality.

This chapter contains the following sections:

- “Considerations for Using the FlowCollector and the Analyzer” section on page 3-2
- “Starting the Analyzer” section on page 3-2
- “Using Data Aggregation Schemes” section on page 3-2
- “File Menu Options” section on page 3-11
- “Edit Menu Options” section on page 3-18
- “Action Menu Options” section on page 3-33
- “Graph Menu Options” section on page 3-45
- “Tools Menu Options for Data Exploration” section on page 3-50
  - “Drilling Down on Network Flows” section on page 3-50
  - “Searching for Flows by Source and Destination Addresses” section on page 3-57
- “Tools Menu Options for Data Collection” section on page 3-65
  - “Controlling TMS Data Collections” section on page 3-65
  - “Controlling NetFlow Data Collections” section on page 3-72
  - “Configuring Routers for Data Export” section on page 3-103
- “Hints Menu Option” section on page 3-111
- “Help Menu Options” section on page 3-112

Considerations for Using the FlowCollector and the Analyzer

Take the following into account when you use the FlowCollector and the Analyzer:

- Only FlowCollector Release 3.0 supports Version 8 export data and the on-router aggregation schemes (see next bullet).
- If you intend to use Analyzer Release 3.0 to process traffic data collected by FlowCollector Release 1.0 or Release 2.0, you cannot get traffic data for the following on-router aggregation schemes (which capture Version 8 export data):
  - RouterAS
  - RouterProtoPort
  - RouterSrcPrefix
  - RouterDstPrefix
  - RouterPrefix
- FlowCollector Release 1.0 does not support the DetailASMatrix aggregation scheme. This scheme was first introduced with FlowCollector Release 2.0 and Analyzer Release 2.0.
- If you intend to use Analyzer Release 3.0 to initiate display functions in the following windows, the DetailASMatrix aggregation scheme must be in effect for the exporting device whose traffic data you wish to view:
  - AS Drill Down window (see the "Tools Menu Options for Data Exploration" section on page 3-50).
  - Search Window (see the "Searching for Flows by Source and Destination Addresses" section on page 3-57).

Starting the Analyzer

For instructions on starting all the Analyzer modules, including the Display module, see the "Installing and Setting Up the Analyzer" section on page 2-1 in Chapter 2.

Using Data Aggregation Schemes

As noted in Chapter 1, the Analyzer displays two types of traffic information:

- NetFlow data
- TMS data

The traffic information for these two data types is processed and arranged into a spreadsheet-like matrix of columns and rows in the display pane of the Display module.

The following sections describe the NetFlow and TMS data aggregation schemes available for use with the FlowCollector and the Analyzer.

NetFlow Data Aggregation Schemes

NetFlow data aggregation schemes consist of key columns and value columns.

Table 3-1 summarizes the NetFlow data aggregation schemes available in the current releases of the FlowCollector and the Analyzer by listing: a) the name of the NetFlow data aggregation scheme; b) the name(s) of the *key* columns incorporated into the Analyzer display output for each aggregation scheme; and c) a description of the contents of each key column in the aggregation scheme output.

Table 3-1 Key Columns of NetFlow Data Aggregation Schemes

Aggregation Scheme Name	Key Column Name in Data Array	Column Contents
SourceNode	Source	IP address of the device that originated the traffic flow.
DestNode	Destination	IP address of the device that received the traffic flow.
HostMatrix	Source Destination	IP address of the source device that originated the traffic flow. IP address of the destination device that received the traffic flow.
Protocol	Protocol	IP transport protocol used in transmitting the monitored traffic. The FlowCollector's nfknown.protocols file defines the recognized application layer protocols (FTP, Telnet, UDP, and so forth) that the FlowCollector recognizes in aggregating Netflow data.
DetailDestNode	Destination SrcPort DestPort Protocol	IP address of the destination device. Application port number of the source device. Application port number of the destination device. IP transport protocol used in transmitting the monitored traffic.
DetailSourceNode	Source SrcPort DstPort Protocol	IP address of the source device. Application port number of the source device. Application port number of the destination device. IP transport protocol used in transmitting the monitored traffic.
DetailHostMatrix	Source Destination SrcPort DestPort Protocol	IP address of the source device. IP address of the destination device. Application port number of the source device. Application port number of the destination device. IP transport protocol used in transmitting the monitored traffic.
DetailInterface	Source Destination Input Output NextHop	IP address of the source device. IP address of the destination device. Input interface number of the device receiving the traffic flow. Output interface number of the device originating the traffic flow. IP address of the next hop device.
SourcePort	SrcPort	Application port number of the source device.
DestPort	DestPort	Application port number of the destination device.
ASMatrix[1]	Source AS Dest AS	Source AS (autonomous system) number. Destination AS number.

Table 3-1 Key Columns of NetFlow Data Aggregation Schemes (continued)

Aggregation Scheme Name	Key Column Name in Data Array	Column Contents
DetailASMatrix[1]	Source	IP address of the source device.
	Destination	IP address of the destination device.
	Source AS	Source AS number.
	Source AS Name	Source AS name.
	Dest AS	Destination AS number.
	Dest AS Name	Destination AS name.
	Input	Input interface number of the device receiving the traffic flow.
	Output	Output interface number of the device originating the traffic flow.
	SrcPort	Application port number of the source device.
	DestPort	Application port number of the destination device.
	Protocol	IP transport protocol used in transmitting the monitored traffic.
NetMatrix[1]	Source	IP address of the source device.
	(Source) Mask	Number of significant bits in the source subnet mask.
	Input	Input interface number of the device receiving the traffic flow.
	Destination	IP address of the destination device.
	(Destination) Mask	Number of significant bits in the destination subnet mask.
	Output	Output interface number of the device originating the traffic flow.
CallRecord	Source	IP address of the source device.
	Destination	IP address of the destination device.
	IP Precedence	**Note:** The six standard "values" columns are not included in the output display for the CallRecord aggregation scheme. Instead, the values columns of this aggregation scheme are:
	TOS	• Active Time—Duration (hold time) of all calls, derived from the elapsed time between the first packet and the last packet in the monitored traffic flow. • Records—Total number of records in the traffic flow. • Packets—Total number of packets in the traffic flow. • Bytes—Total number of bytes in the traffic flow. • Flows—Total number of traffic flows monitored.
RouterAS[2]	Source AS	Autonomous system (AS) from which the monitored traffic originated.
	Source AS Name	Name of the source AS system.
	Dest AS	AS system name to which the monitored traffic was delivered.
	Dest AS Name	Name of the destination AS system.
	Input Interface	On the router from which the data was collected, the ifIndex number and description text for the physical interface through which the monitored traffic was received.
	Output Interface	On the router from which the data was collected, the ifIndex number and description text for the physical interface through which the monitored traffic was sent.
RouterProtoPort[2]	Source Port	On source host, the application port from which the monitored traffic was sent. This may be a port number or (if the FlowCollector is so configured) a text string.
	Destination Port	On destination host, the application port number to which the monitored traffic is delivered. This may be a port number or a text string.
	Protocol	IP protocol used in transmitting the monitored traffic.

Table 3-1 Key Columns of NetFlow Data Aggregation Schemes (continued)

Aggregation Scheme Name	Key Column Name in Data Array	Column Contents
RouterPrefix[2]	Source (Subnet)	IP address of the network from which the monitored traffic originated.
	Destination (Subnet)	IP address of the network to which the monitored traffic was delivered.
	Source (Prefix) Mask	Mask by which the source IP address was subnetted.
	Destination (Prefix) Mask	Mask by which the destination IP address was subnetted.
	Input Interface	On the router from which the data was collected, the ifIndex number and description text for the physical interface through which the monitored traffic was received.
	Output Interface	On the router from which the data was collected, the ifIndex number and description text for the physical interface through which the monitored traffic was sent.
	Source AS	Source autonomous system (AS) number.
	Source AS Name	Name of source AS.
	Destination AS	Destination AS number.
	Destination AS Name	Name of destination AS.
RouterDstPrefix[2]	Destination (Subnet)	IP address of the network to which the monitored traffic was delivered.
	Destination (Prefix) Mask	Mask by which the destination IP address was subnetted.
	Output Interface	On the router from which the data was collected, the ifIndex number and description text for the physical interface through which the monitored traffic was sent.
	Destination AS	Destination AS number.
	Destination AS Name	Name of destination AS.
RouterSrcPrefix[2]	Source (Subnet)	IP address of the network from which the monitored traffic originated.
	Source (Prefix) Mask	Mask by which the source IP address was subnetted.
	Input Interface	On the router from which the data was collected, the ifIndex number and description text for the physical interface through which the monitored traffic was received.
	Source AS	Source AS number.
	Source AS Name	Source AS name.

1 Compatible only with Version 5 and Version 7 export data (Release 2.0 or later of the FlowCollector).
2 Compatible only with Release 3.0 of the FlowCollector and Version 8 export data (for on-router data aggregation).

The *key* columns of a NetFlow data aggregation scheme represent the traffic information that the FlowCollector looks for in sorting and processing UDP datagrams received from NetFlow exporting devices in your network. Once the FlowCollector processes the data according to the aggregation scheme currently in effect, it stores the data in its local NetFlow directory.

The *value* columns of a NetFlow data aggregation scheme contain the statistical information extracted from the UDP datagrams for a given traffic flow. The value columns, which are common to all of the NetFlow data aggregation schemes (except for the CallRecord aggregation scheme, as noted in Table 3-1), are listed and described in Table 3-2.

For more detailed information about the FlowCollector and the associated NetFlow data aggregation schemes, consult the *NetFlow FlowCollector Installation and User Guide*.

Table 3-2 Value Columns of NetFlow Data Aggregation Schemes

Value Column Name	Value Column Contents
Packets	Total number of packets in the monitored traffic flow.
Pkts/Sec	Packet rate per second of the monitored traffic flow.
Bytes	Total number of bytes in the monitored traffic flow.
Kbits/Sec	Total bit rate in kilobits per second.
Flows	Total number of monitored traffic flows.
Flows/Min	Number of flows monitored per minute.

Key Columns of NetFlow Data Aggregation Schemes

Figure 3-1 shows a sample display of a typical NetFlow data aggregation scheme, in this case, the DetailInterface aggregation scheme.

Due to space limitations in the display pane, not all of the columns in an aggregation scheme can be displayed in the pane at one time. Hence, aggregation scheme data typically overflows the available space in the display pane. Therefore, vertical and horizontal scroll bars are incorporated into the pane to enable you to traverse to any area of interest in the overall data array.

Figure 3-1 Key Columns of a Typical NetFlow Data Aggregation Scheme

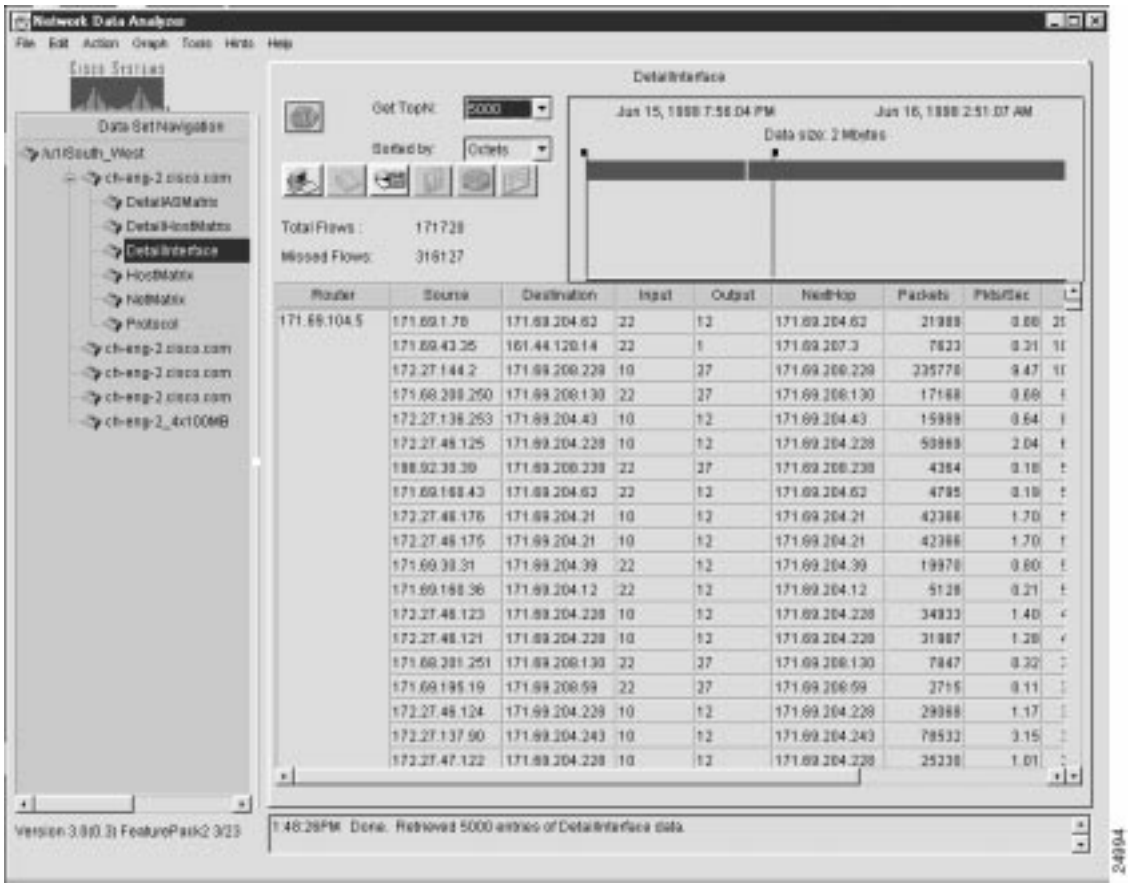

The Router column shown in Figure 3-1 is common to all of the NetFlow data aggregation schemes. This column identifies the source device from which the displayed NetFlow data was collected.

The *key* columns of a displayed NetFlow data aggregation scheme contain information peculiar to the particular scheme being displayed.

In Figure 3-1, for example, the key columns of the DetailInterface aggregation scheme include the following:

- Source
- Destination
- Input
- Output
- NextHop

These columns identify the specific subset of sort keys used by the FlowCollector in post-processing the UDP export datagrams received from the router identified (by the IP address) in the first column of the display.

Using these sort keys, the FlowCollector filters and aggregates the appropriate data from the exported UDP datagrams and stores the results as a specific collection of data in its local NetFlow repository.

- Values columns—The Packets column marks the beginning of the values columns in the displayed aggregation scheme (see the next section). The values columns are always the six rightmost columns of a displayed NetFlow data aggregation scheme.

Value Columns of NetFlow Data Aggregation Schemes

Figure 3-2 is a representation of the same sample output for the DetailInterface aggregation scheme shown in Figure 3-1, except that the display pane is scrolled to the right to make all of the *value* columns in the data array visible.

The value columns of a displayed NetFlow data aggregation scheme always comprise the six rightmost columns of the data array, as identified below:

- Packets
- Pkts/Sec
- Bytes
- Kbits/sec
- Flows
- Flow/Min

Figure 3-2 Value Columns of Typical NetFlow Data Aggregation Scheme

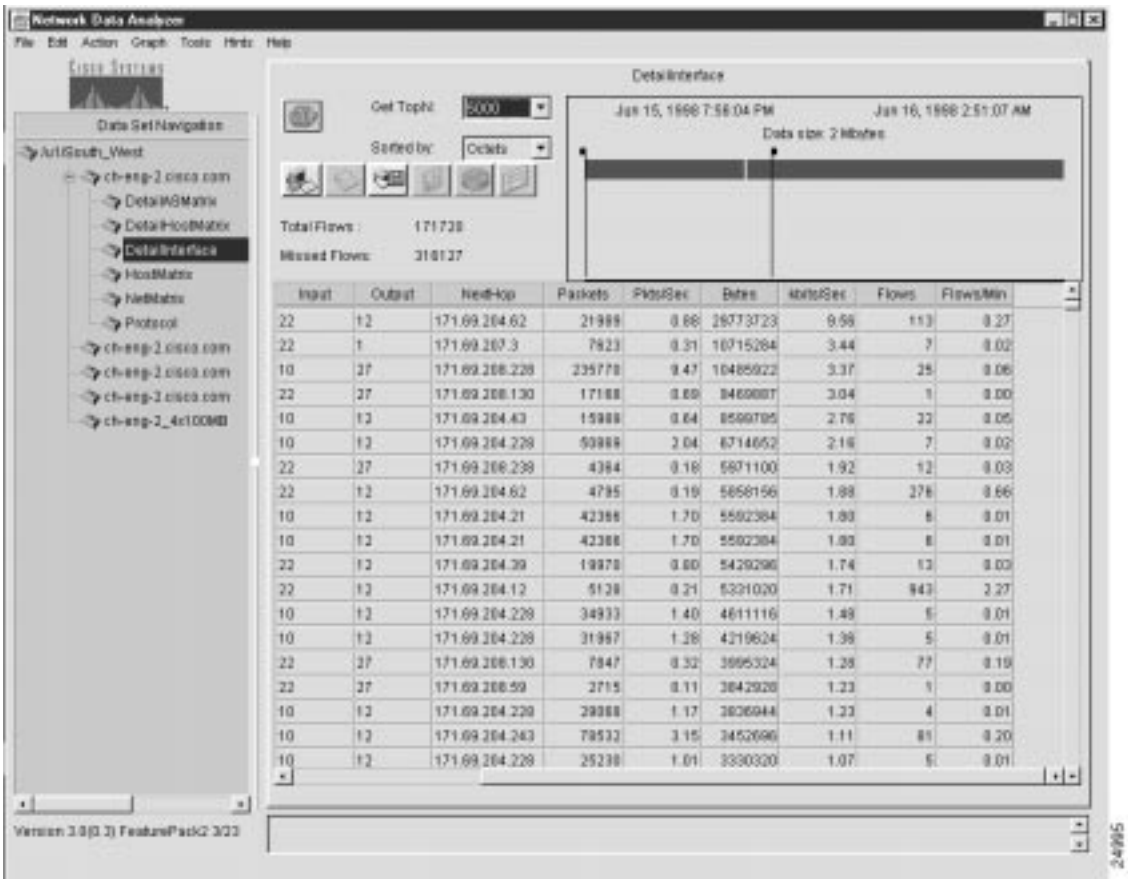

TMS TrafficMatrix Data Aggregation Scheme

Table 3-3 describes the two key columns of the TMS TrafficMatrix data aggregation scheme, while Table 3-4 describes the value columns of the TMS TrafficMatrix data aggregation scheme.

Note that the value columns in Table 3-4 pertain to two different types of traffic: E (external) packets, and I (internal) packets.

Table 3-3 Key Columns of TMS TrafficMatrix Data Aggregation Scheme

Key Field Name	Description
Destination	The destination IP address and destination IP address mask for type “p” records. Type “p” records are indexed by destination prefix and describe dynamic tag switching traffic data or traffic engineered (TE) tunnel head traffic data.
Tunnel ID	The tunnel head IP address and tunnel serial ID for type “t” records. Type “t” records are indexed by tunnel head for tunnel midpoint records and describe traffic engineered (TE) tunnel midpoint data.

Table 3-4 Value Columns of TMS TrafficMatrix Data Aggregation Scheme

Name of Value Field	Description
E.pkts	External packet count within the specified start/stop interval for the flow.
E.Pkts/Sec	External packet rate, in packets per second.
E.Bytes	External byte count within the specified start/stop interval for the flow.
E.kbits/Sec	External byte rate, in kilobits per second.
I.Pkts	Internal packet count within the specified start/stop interval for the flow.
I.Pkts/Sec	Internal packet rate, in packets per second.
I.Bytes	Internal byte count within the specified start/stop interval for the flow.
I.kbits/Sec	Internal byte rate, in kilobits per second.
Pkts	Total packet count within the specified start/stop interval for the flow.
Pkts/Sec	Total packet rate, in packets per second.
Bytes	Total byte count within the specified start/stop interval for the flow.
kbits/Sec	Total bit rate, in kilobits per second.
Route Flaps	Number of route flaps detected within the specified start/stop interval for the flow. The term route flap refers to an instance of the route going down and being restored.
Flaps/Min	Measured rate of route flaps per minute within the specified start/stop interval for the flow.

Key Columns of TMS TrafficMatrix Data Aggregation Scheme

Figure 3-3 shows sample output for the TMS TrafficMatrix data aggregation scheme, which is designed specifically for use in collecting and displaying TMS data.

Figure 3-3 Key Columns of TMS TrafficMatrix Data Aggregation Scheme

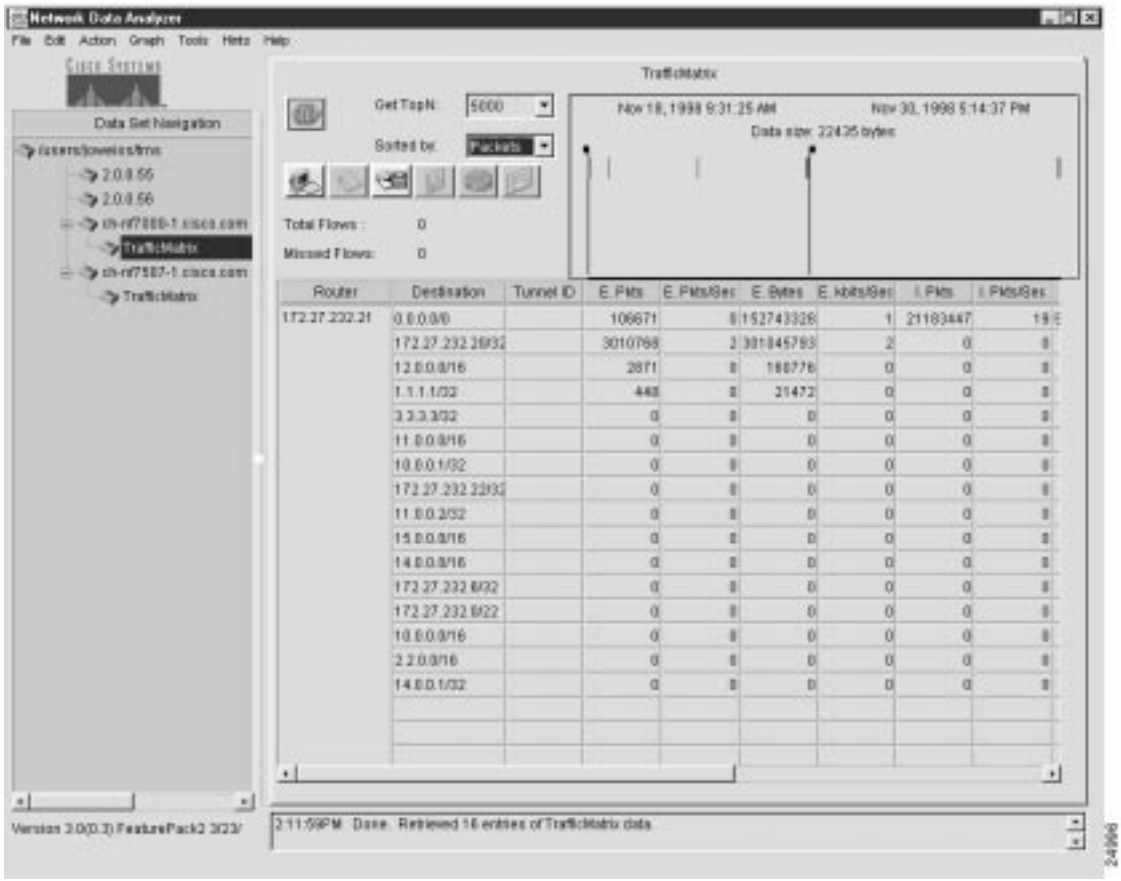

The Router column of the TMS data array (see Figure 3-3) always identifies the source device from which the displayed TMS data was collected.

The *key* columns of a TMS data array include the following:

- Destination column
- Tunnel ID column
- Values columns—The E.Pkts column marks the beginning of the values columns in the displayed TMS aggregation scheme. These values columns (see Table 3-4) are always the fourteen rightmost columns of a displayed TMS data aggregation scheme.

Value Columns of TMS TrafficMatrix Data Aggregation Scheme

Figure 3-4 shows the same sample output pertaining to the TrafficMatrix aggregation scheme as that shown in Figure 3-3, except that the data in the display pane is shifted to the left by means of the horizontal scroll bar, bringing most of the value columns comprising the overall data array into view.

The E.Pkts column marks the beginning of the 14 *values* columns displayed in a TMS aggregation scheme. Table 3-4 describes the 14 different values columns that appear in a TMS TrafficMatrix aggregation scheme.

Figure 3-4 Value Columns of TMS TrafficMatrix Data Aggregation Scheme

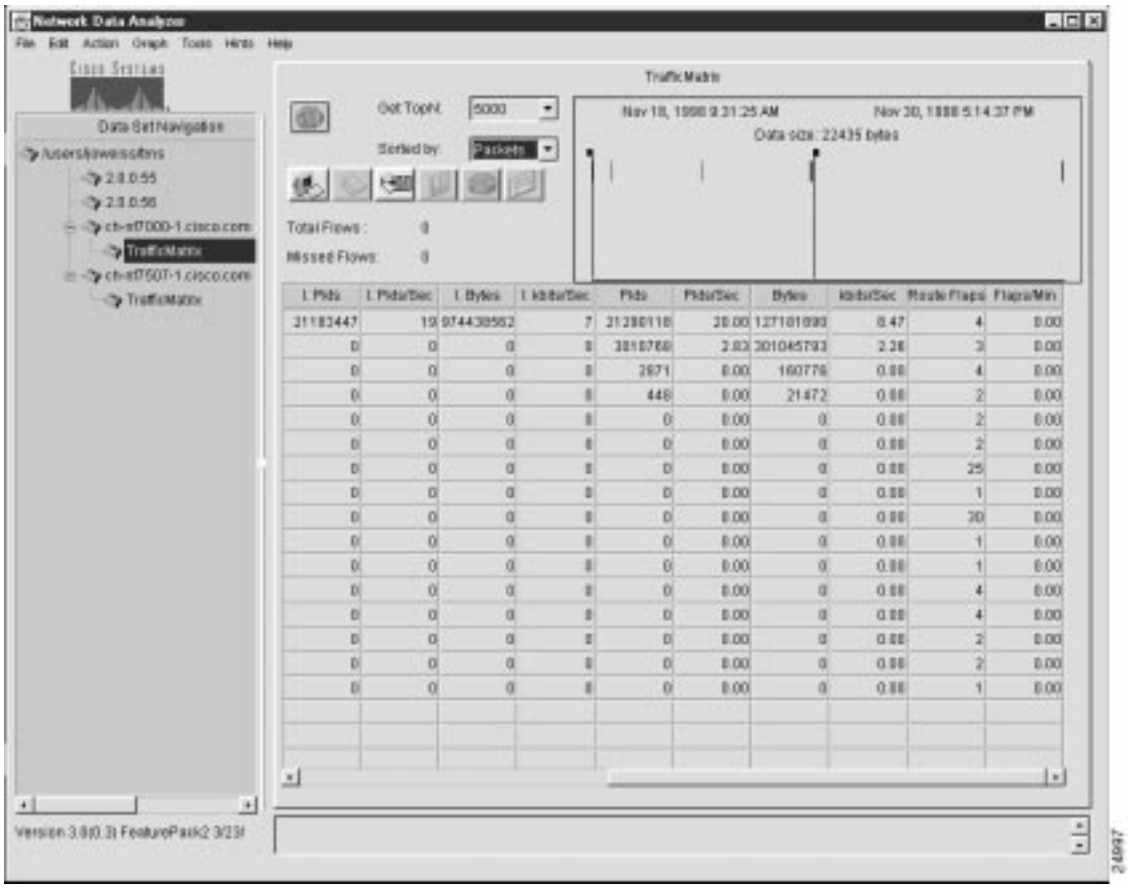

File Menu Options

The Display module File menu provides the following selectable options:

- Export—See the “Exporting Aggregation Scheme Data to a Named File” section on page 3-12.
- Load—See the “Loading a Tree File” section on page 3-15.
- Save—See the “Creating and Saving a Tree File” section on page 3-13.
- Quit—See the “Loading a Tree File” section on page 3-15.

The menu system provided by the Analyzer is structured to help you accomplish specific tasks in a logical fashion, generally proceeding from simple to more complex tasks.

Some menu options are grayed out (not selectable) until other requisite conditions are satisfied.

The following sections describe how you use the File menu options.

Exporting Aggregation Scheme Data to a Named File

You can export (save) the contents of any displayed aggregation scheme to a named file at any time. Figure 3-5 shows a typical "DetailInterface" aggregation scheme that you could use as the basis for the Export function of the File menu.

Figure 3-5 Sample Data Aggregation Scheme for File Export Operations

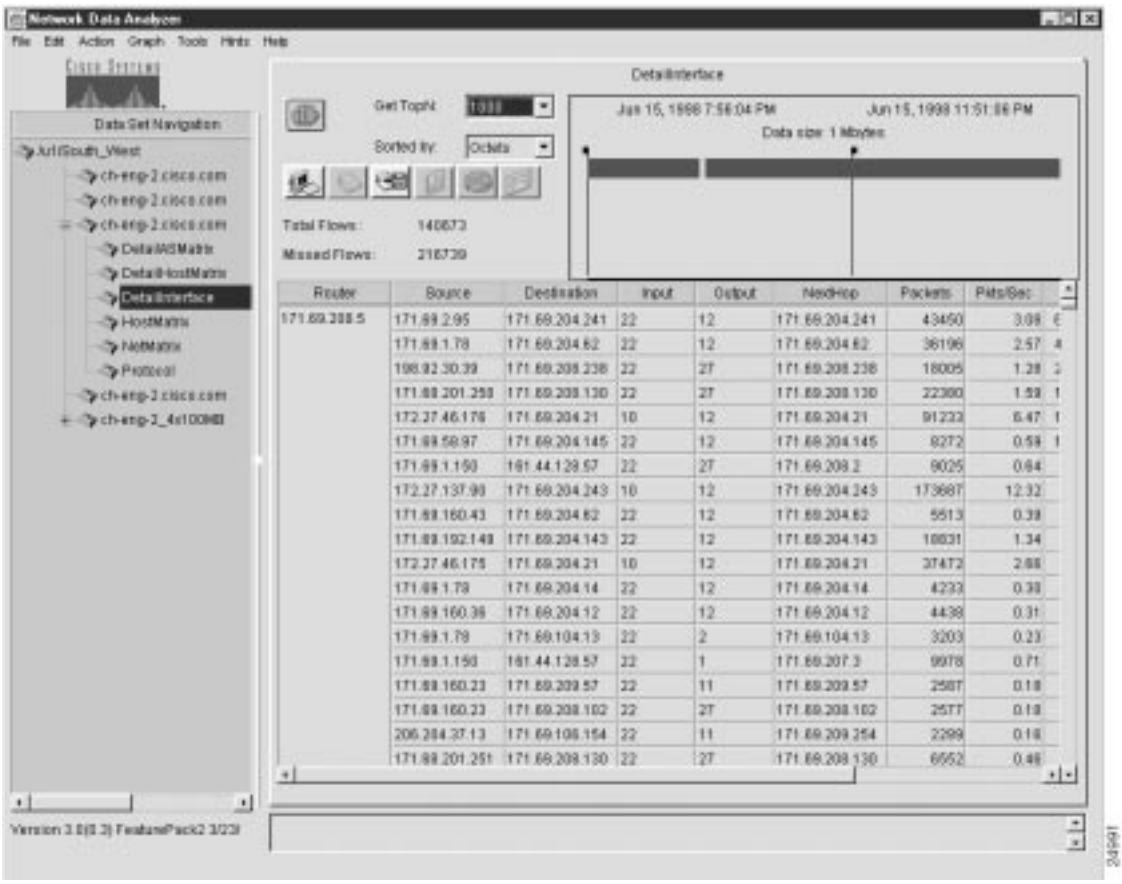

To export the contents of a displayed aggregation scheme to a named file, perform the following procedure:

Step 1 If the data aggregation scheme that you want to export is currently displayed (as in Figure 3-5), select the Export option of the File menu or click the Export Data button in the display pane.

Either action causes the following pop-up window to appear.

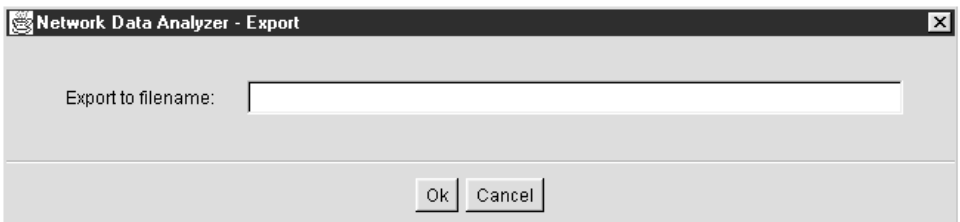

Step 2 Choose a name by which you want the displayed aggregation scheme data to be saved, such as “detail_interface_output.”

Step 3 Enter the name of the file in the space provided in the export pop-up window, as shown below.

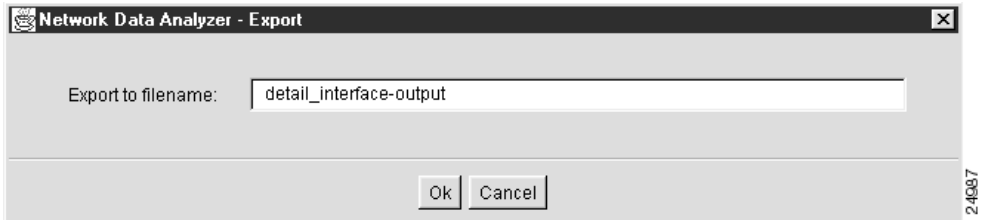

Step 4 Click OK.

The following message box appears, identifying the saved file by name and indicating the location in which it is stored.

The export file is stored in comma separated vector (.CSV) format in a specific directory on the device presently serving as the DisplayServer host (fredm-ultra, for purposes of this procedure). The .CSV file format can be interpreted by Microsoft Excel and other popular spreadsheet programs.

You can save aggregation scheme data as a .CSV file and import the file into any program that can handle this file format.

Creating and Saving a Tree File

Operational conveniences and efficiencies can be realized by creating and saving named tree files. After you create a tree file, you can load it into the Data Set Navigation pane at any time and use it as the basis for a current Display module session.

In using the Analyzer to meet day-to-day needs, it is to your advantage to create a library of tree files that you can draw upon. For example, you can compose a working data tree structure that encompasses exporting devices in a certain network segment or geographical area of interest.

With a library of tree files available when you start up the Display module, you can load any desired working data tree structure into the Data Set Navigation pane as a single entity, thus avoiding having to populate the Data Set Navigation pane individually with data set paths in preparation for a given Display module session.

After you populate the Data Set Navigation pane with the desired data set paths, save the contents of the pane as a named file by invoking the Save option of the File menu.

To create and save a tree file, perform the following steps:

Step 1 Add a desired number of data set paths one at a time into the Data Set Navigation pane.

The procedure for populating the Data Set Navigation pane with desired data set paths is described in the "Adding Data Set Paths to the Data Set Navigation Pane" section on page 3-18.

For purposes of creating and saving a named tree file, assume that you have added the data set paths shown in Figure 3-6 into the Data Set Navigation pane.

Figure 3-6 Content of Data Set Navigation Pane for Creating a Named Tree File

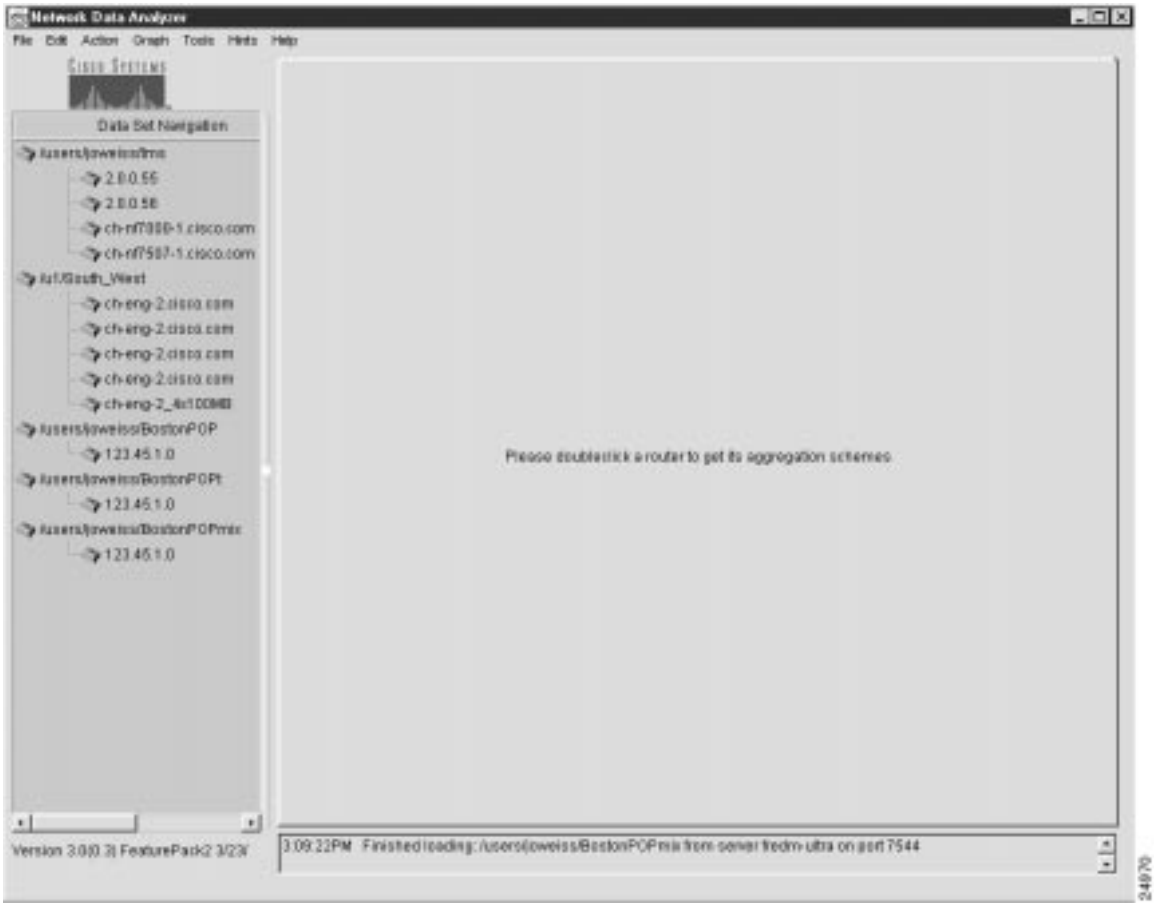

Step 2 After populating the Data Set Navigation pane with the desired data set paths, save the data sets paths as a tree file by invoking the Save option of the File menu.

This action causes the following Configurations pop-up window to appear.

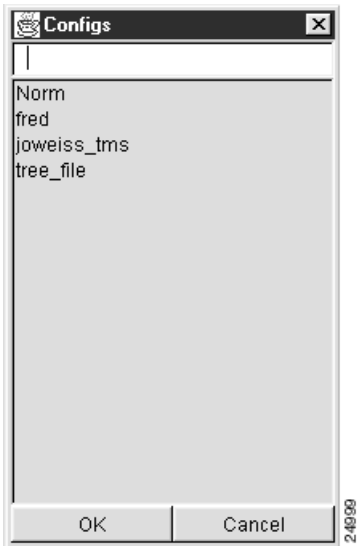

Step 3 In the data field at the top of the Configurations pop-up window, enter a file name of your choosing, such as “Boston_Pops,” by which the tree file is to be identified.

Step 4 Click OK to save the file by the designated name in a directory on the DisplayServer module host.

Loading a Tree File

You can select the Load option from the File menu at any time to select any one of several pre-defined tree files for loading into the Data Set Navigation pane.

The “Creating and Saving a Tree File” section on page 3-13 describes how to create a library of pre-defined tree files that you can load into the Data Set Navigation pane to serve as the selection mechanism for Display module functions.

To load a tree file into the Data Set Navigation pane, perform the following procedure:

Step 1 Select the Load option from the File menu. The following pop-up window appears:

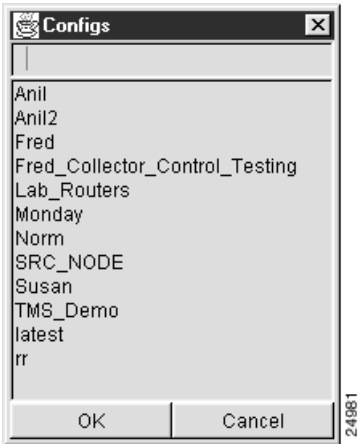

As evident from the contents of this pop-up window, several tree files have been created and saved previously for use in performing current or intended Display module tasks.

Step 2 From those tree files listed in the Configurations pop-up window, determine which file you want to load into the Data Set Navigation pane.

For this purpose, assume that you want to load the tree file named “Norm.”

Step 3 To load the selected tree file, do either of the following:

- Click the file to highlight it, causing its name to appear in the box at the top of the pop-up window, as shown below.

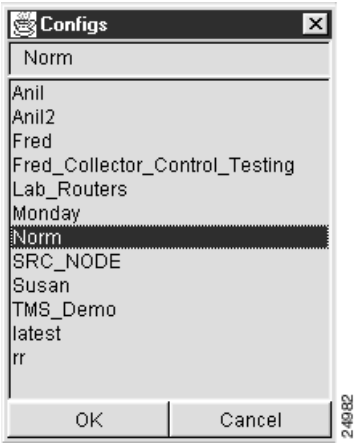

  Click OK at the bottom of the pop-up window.

- Double-click the name of the selected file.

Either action above causes the selected tree file to be loaded into the Data Set Navigation pane (see Figure 3-7).

Figure 3-7 Selected Tree File Loaded into the Data Set Navigation Pane

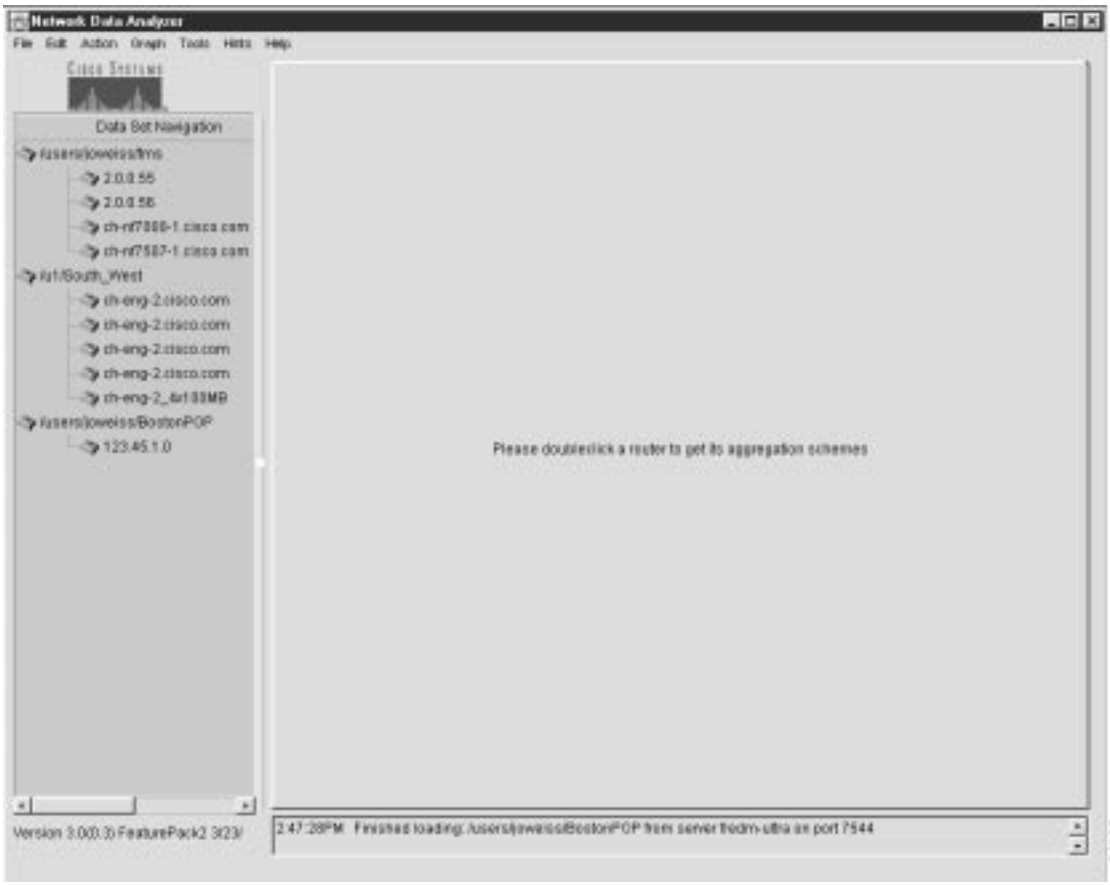

With the data tree structure shown in Figure 3-7 loaded into the Data Set Navigation pane, you can invoke any Display module function, provided that the selected option is not grayed out (unavailable in the current operating context of the Analyzer).

The following loading options become available to you once you have created and saved one or more tree files:

- Option 1—You can load any tree file by means of the procedure outlined above.

  Typically, on Display module startup, you would load the contents of a tree file to establish the working data tree structure for the current Display module session.

  During this same session, you can load any other desired tree file, as needed, in which case, the current working data tree structure is overwritten. Note, however, that the previous tree structure is still preserved for future use.

  Tree files are stored on a DisplayServer host according to the parameters that you specify in the /opt/CSCOnfa/NFADisplay/bin/start.Display file.

- Option 2—You can edit the Display module's start.Display file to cause a default tree file to be loaded automatically by default into the Data Set Navigation pane on Display module startup.

  In this case, you must include the name of a user-defined tree file as an argument to the defaulttreefile keyword in the start.Display file. After you reference a default tree file in the start.Display file, this file takes precedence in populating the Data Set Navigation pane with a working data tree structure on Display module startup.

In effect, the start.Display file serves as a startup script for each new Display module session that you initiate from scratch. Thus, you have the ability to establish beforehand the working data tree structure that will be loaded automatically into the Data Set Navigation pane on Display module startup.

The "Optional Setup Procedures" section on page 2-4 in Chapter 2 provides instructions for editing the start.Display file. If you do not create a default tree file and include its name as an argument to the `defaulttreefile` keyword in the start.Display file, the Data Set Navigation pane will be blank on Display module startup.

The default data tree structure loaded into the Data Set Navigation pane on Display module startup remains in effect until you deliberately change it, delete it, or replace it.

When you want to work with a data tree structure other than the one currently being displayed in the Data Set Navigation pane, you can use the Load option of the File menu, as described above, to load any other desired tree file.

Exiting from the Analyzer

To exit from the Analyzer, select the Quit option of the File menu.

Edit Menu Options

This section describes how to use the following options on the Edit menu:

- Add Data Path—See the "Adding Data Set Paths to the Data Set Navigation Pane" section on page 3-18.
- Add Router Group—See the "Creating and Saving a Named Router Group" section on page 3-20.
- Remove Router or Group—See the "Removing a Data Set Path or a Named Router Group" section on page 3-25.
- Properties—See the "Displaying the Properties of Data Set Paths" section on page 3-28.

Adding Data Set Paths to the Data Set Navigation Pane

This section describes how to manually add individual data set paths to the Data Set Navigation pane.

To group individual data set paths into named router groups that appear automatically in the Data Set Navigation pane, refer to the "Creating and Saving a Named Router Group" section on page 3-20.

To retrieve and analyze either NetFlow data or TMS traffic data on initial startup of the Display module, you must know beforehand the location of specific directories in a FlowCollector host or an NFS-mounted storage volume in which desired traffic information is known to be stored. Hence, you must add to the Data Set Navigation pane those data set paths that point to NetFlow or TMS traffic data of interest.

To add individual data set paths manually to the Data Set Navigation pane, perform the following steps:

Step 1 From the Display module pull-down menu, select the Add Data Path option of the Edit menu. The New Data Location dialog box appears.

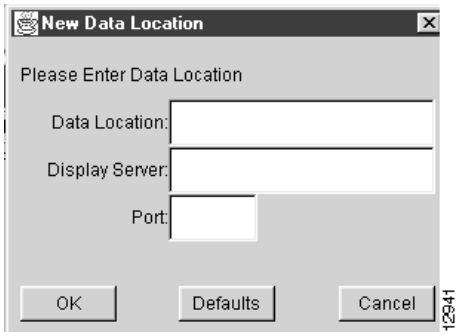

Step 2 Enter the appropriate information to define the desired data set path:

(a) Data Location—This field identifies the directory in a storage repository (/u1/South_West, for example) that you know contains the NetFlow desired data.

You must uniquely define this directory so that it can be accessed directly by the DisplayServer module specified in (b) below.

(b) Display Server—This field identifies the DisplayServer host (fredm-ultra, for example) that has been configured to service user requests for traffic data. The DisplayServer retrieves the requested data from the directory specified in (a) above.

You can identify the DisplayServer host by either its logical name (fredm-ultra, for example) or its IP address.

(c) Port—This field identifies the application port number used by the DisplayServer specified in (b) above in servicing user data requests.

Typically, application port number 7544 is used by the DisplayServer in listening for user commands.

Step 3 When you complete the fields of the New Data Location dialog box, as shown below, click OK.

This step adds the specified new data set path into the Data Set Navigation pane, as shown in Figure 3-8.

By repeating this procedure, you can add multiple data set paths to the Data Set Navigation pane.

To conserve available space in the Data Set Navigation pane, add only those data set paths relevant to your current Display module session.

Note When the Version 3.0 FlowCollector is running with the NFC20_COMPATIBLE_MODE flag set to "no" in its nf.resources file, an extra node is added to the directory structure for collections that include the thread ID. This extra node must be taken into account when specifying the location of data collected in this manner.

Figure 3-8 New Data Set Path Added to Data Set Navigation Pane

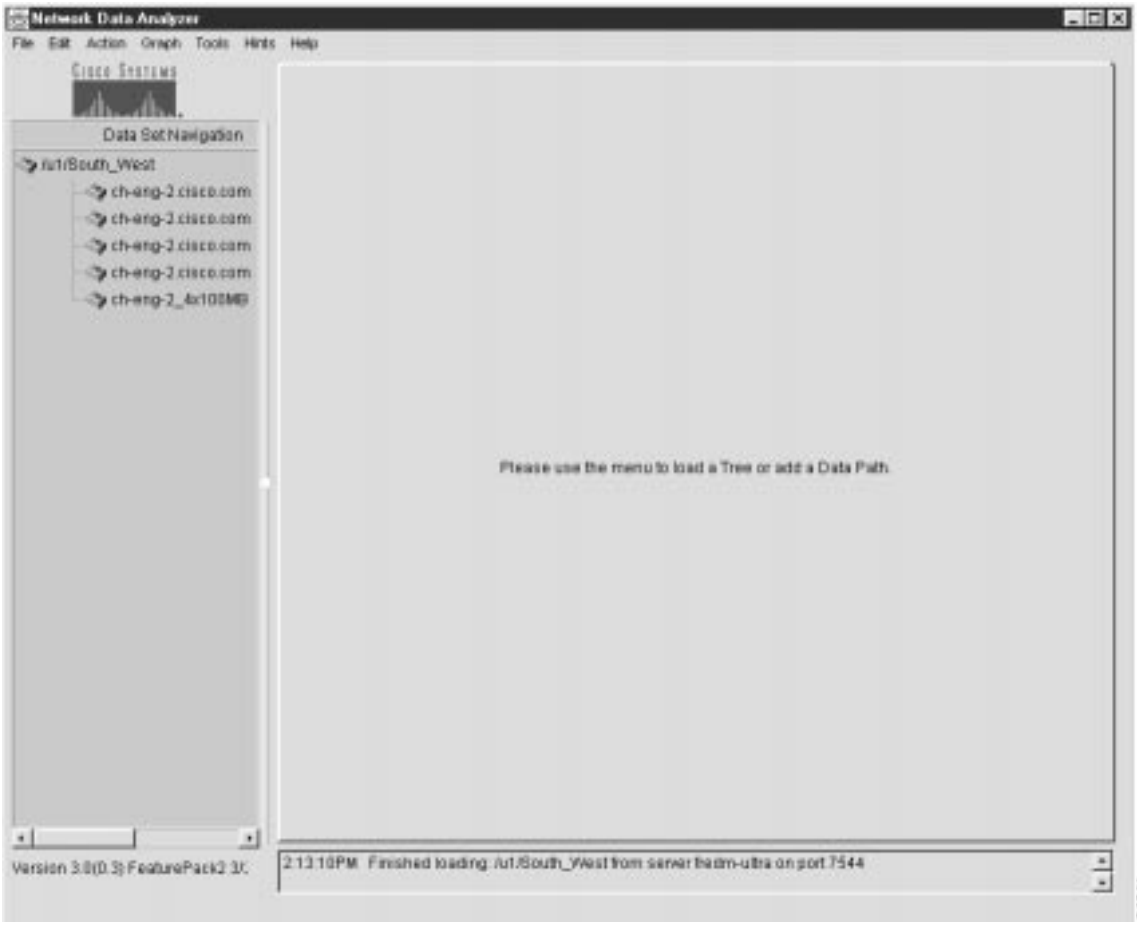

Note Note that the data exporting device "ch-eng-2.cisco.com" is listed multiple times in the Data Set Navigation pane shown in Figure 3-8. This is not an unusual circumstance, since the same device can be exporting NetFlow data from multiple ports.

Creating and Saving a Named Router Group

It can be useful to combine several individual data set paths into a named router group that relates to a particular set of routers or switches in your network.

A named router group constitutes a working data tree structure that you can load into the Data Set Navigation pane at any time as a single entity, obviating the need to populate the Data Set Navigation pane manually with data set paths at each startup of the Display module. Defining a named router group for each of your operating needs enables you to load, change, and delete named router groups at will to serve a variety of Analyzer operating needs.

Note Each data set path can be assigned to one or more named router groups using the "clone" option of the Properties menu. This does not alter the links to directories in the designated NetFlow or TMS data storage facilities. The integrity of all data set paths and associated directory pointers is preserved.

For more information on cloning, see the "Cloning Function of the Properties Window" section on page 3-30.

To group multiple data set paths into a named router group, perform the following steps:

Step 1 Populate the Data Set Navigation pane with any number of data set paths of relevance for defining a named router group.

If the Data Set Navigation pane is not already populated with appropriate data set paths (such as those shown in Figure 3-9) add any number of data set paths individually to the Data Set Navigation pane through successive uses of the Add Data Path option of the Edit menu.

Figure 3-9 Data Tree Structure for Defining a Named Router Group

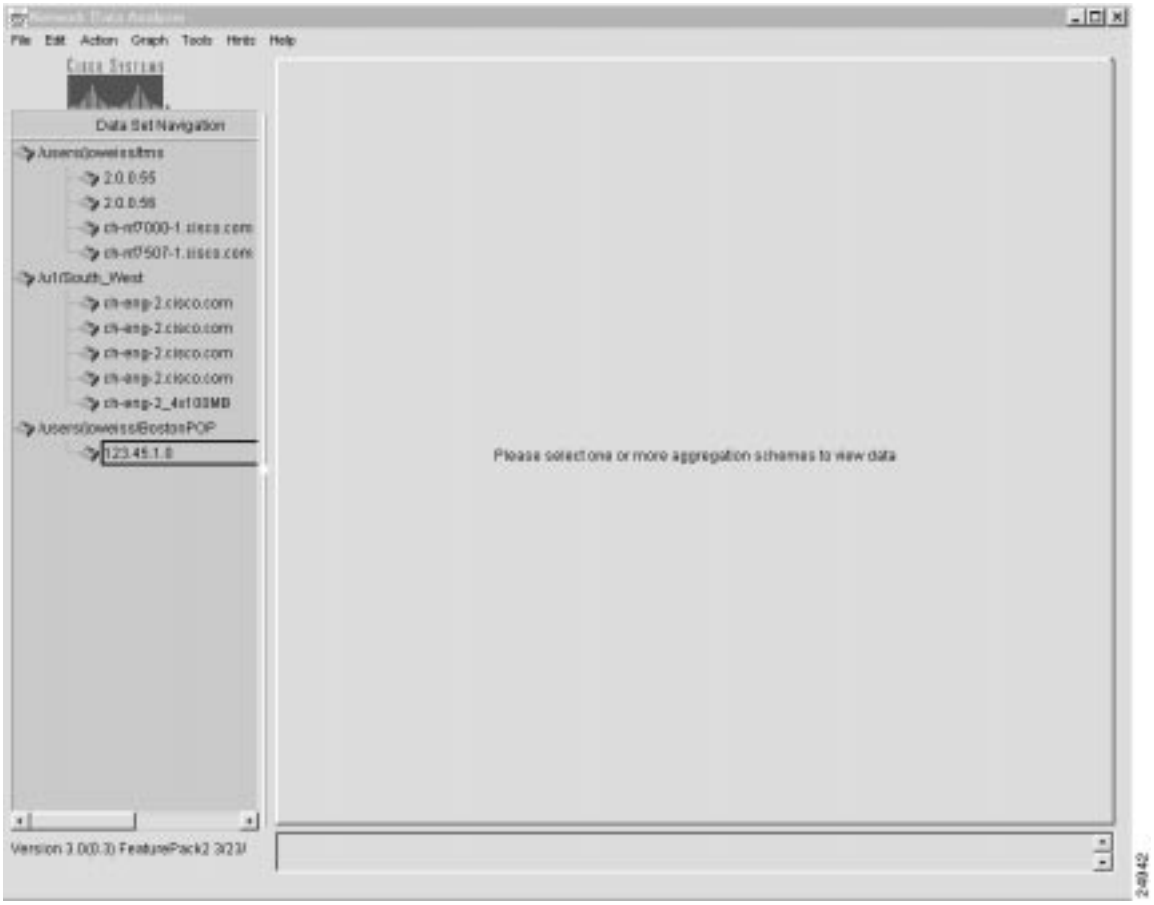

The procedure for this function is described in the "Adding Data Set Paths to the Data Set Navigation Pane" section on page 3-18.

Step 2 Add the Add Router Group name to the Data Set Navigation pane by doing the following:

(a) Select the Add Group option from the Edit menu. The following pop-up window appears.

(b) In this Set Group Name window, enter a router group name of your choice, such as "North_American_Region," for example.

(c) Click OK.

The router group name "North_American_Region" then appears at the bottom of the Data Set Navigation pane (see Figure 3-10).

Figure 3-10 User-Defined Name for Router Group

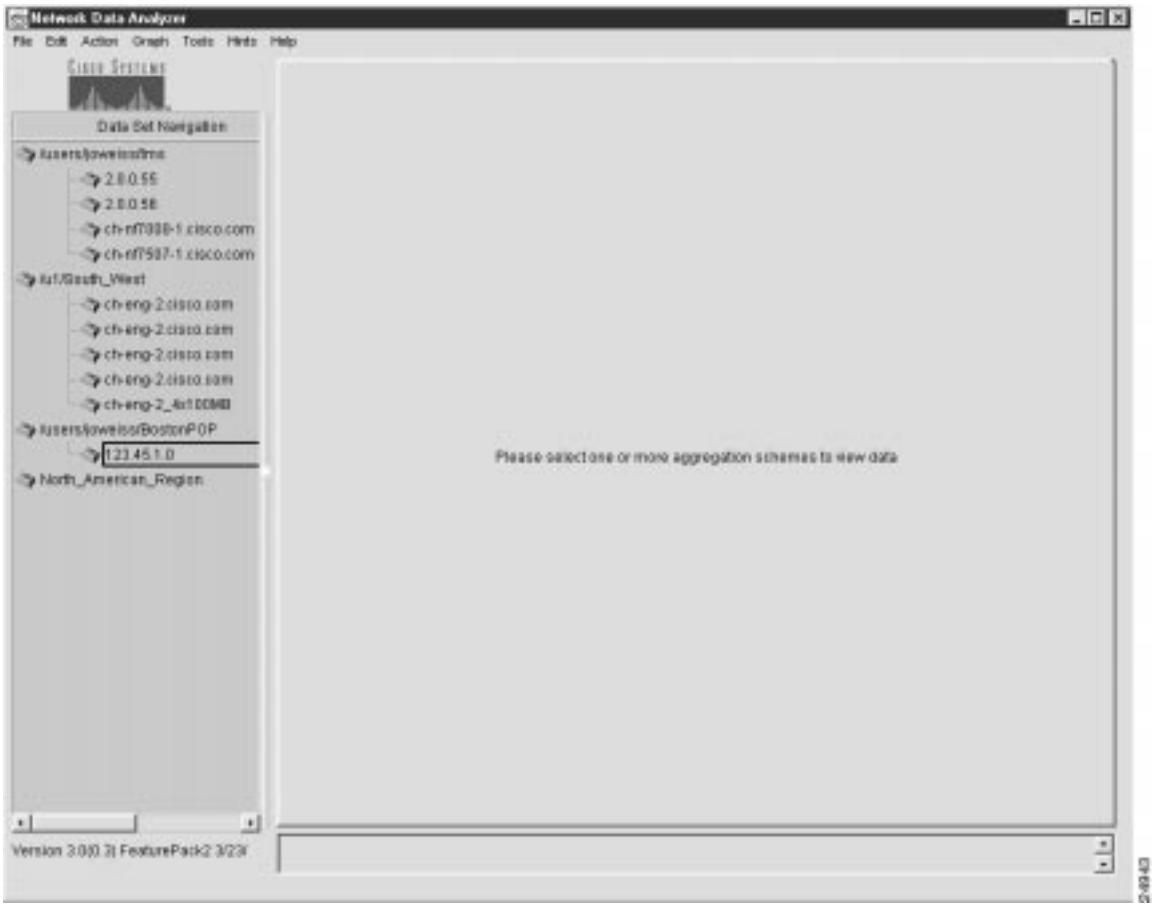

Step 3 Populate the named router group with data set paths.

Using the data set paths already present in the Data Set Navigation pane, drag the desired data set paths one at a time into the router group named "North_American_Region."

Click the desired data set path to highlight it. Hold down the left mouse button and drag the data set path into the named router group.

Repeat this step as many times as necessary to compose the named router group. For this example, assume that you want to drag the data set paths "/users/joweiss/tms," "/u1/South_West," and "/users/joweiss/BostonPOP" into the named router group.

Note Each data set path can be assigned to one or more named router groups using the "clone" option of the Properties menu. This does not alter the links to directories in the designated NetFlow or TMS data storage facilities. The integrity of all data set paths and associated directory pointers is preserved.

Step 4 After composing the named router group, you can:

- Leave all the data set paths in the Data Set Navigation pane intact if you foresee a need for them in the current Display module session.
- Delete the data set paths to conserve space in the Data Set Navigation pane.

If you choose the second option, delete any non-essential data set paths by successively invoking the Remove Router Or Group option of the Edit menu, as described in the next section.

Step 5 Save the existing data tree structure to a file. To do so:

(a) Select the Save option from the File menu. The following pop-up window appears:

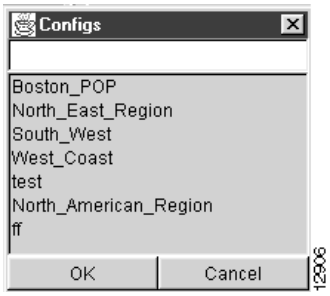

(b) Enter a file name of your choice in the blank field at the top of the window to identify the saved file. Assume that you want to name the file "North_American_Region." The pop-up window appears as follows:

(c) Click OK to save the file.

On completion of this procedure, the Data Set Navigation pane contains the named router group shown in Figure 3-11.

Figure 3-11 Working Tree Structure of a Named Router Group

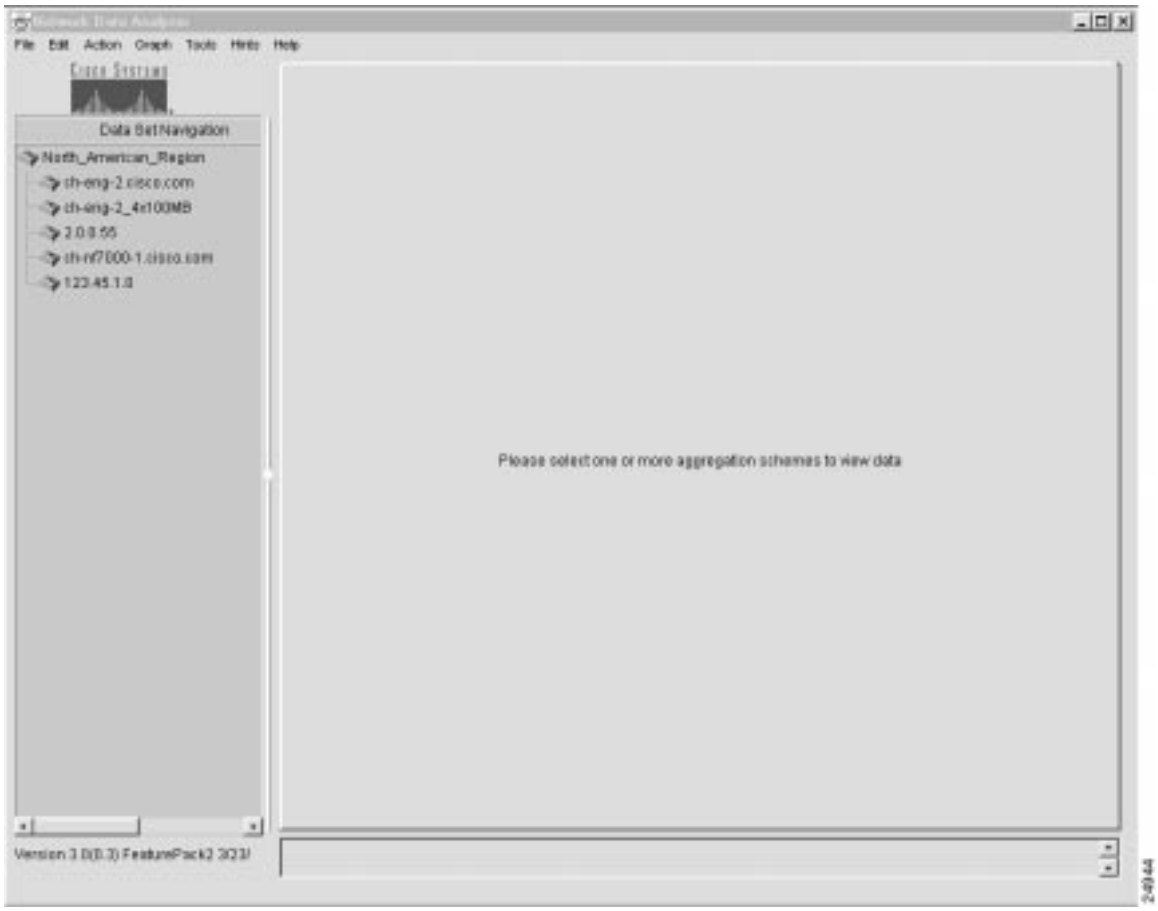

With this working data tree structure in place in the Data Set Navigation pane, you can select any available aggregation scheme in the pane as the basis for Display module tasks.

Removing a Data Set Path or a Named Router Group

This section describes how to “clean up” the existing data tree structure in the Data Set Navigation pane by means of the Remove Router Or Group option of the Edit menu.

Your existing data tree structure might become outdated due to changes in your networking or data exporting environment. Such changes may include:

- Additions to or deletions from the complement of data exporting devices in your network.
- Changes in the number or location of host FlowCollector or NFS-mounted devices in your network for storing exported traffic data.
- Changes in the way you want to group data set paths for operational convenience.

Any fundamental configuration changes to FlowCollectors or the directory pointers for stored traffic information can potentially invalidate an existing data set path or an existing element of a named router group.

For this reason, the Display module incorporates capabilities that enable you to tailor the contents of the Data Set Navigation pane.

The sample tree structure shown in Figure 3-12 shows a context for removing selected elements of a data tree structure from the Data Set Navigation pane.

Figure 3-12 Context for Removing Elements of Data Tree Structure

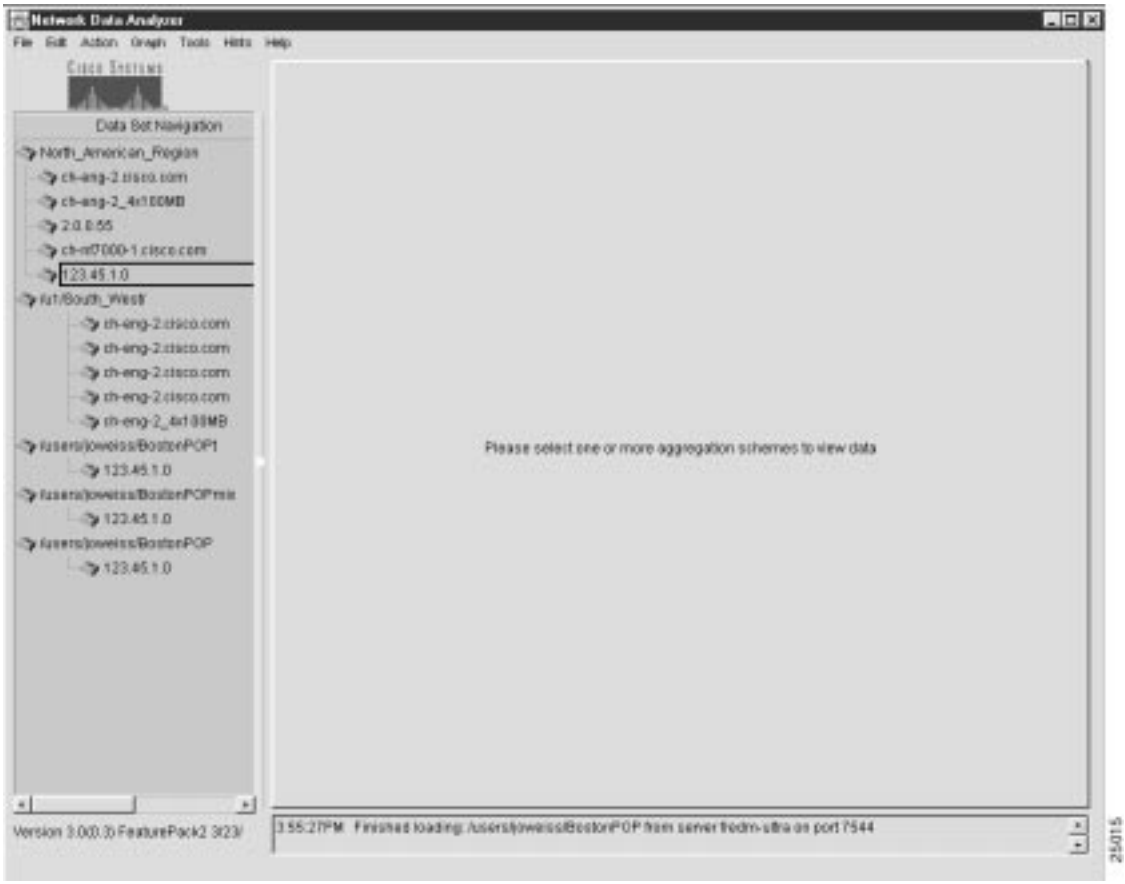

To remove any element of the tree structure currently being displayed in the Data Set Navigation pane, perform the following procedure:

Step 1 Determine which element of an existing data tree structure that you want to remove.

For purposes of this step, assume that you want to delete the entire router group named "u1/South_West."

Step 2 Click the router group name "u1/South_West" to highlight it.

This action changes the appearance of the Display module window (see Figure 3-13.

Figure 3-13 Selecting a Named Router Group for Removal

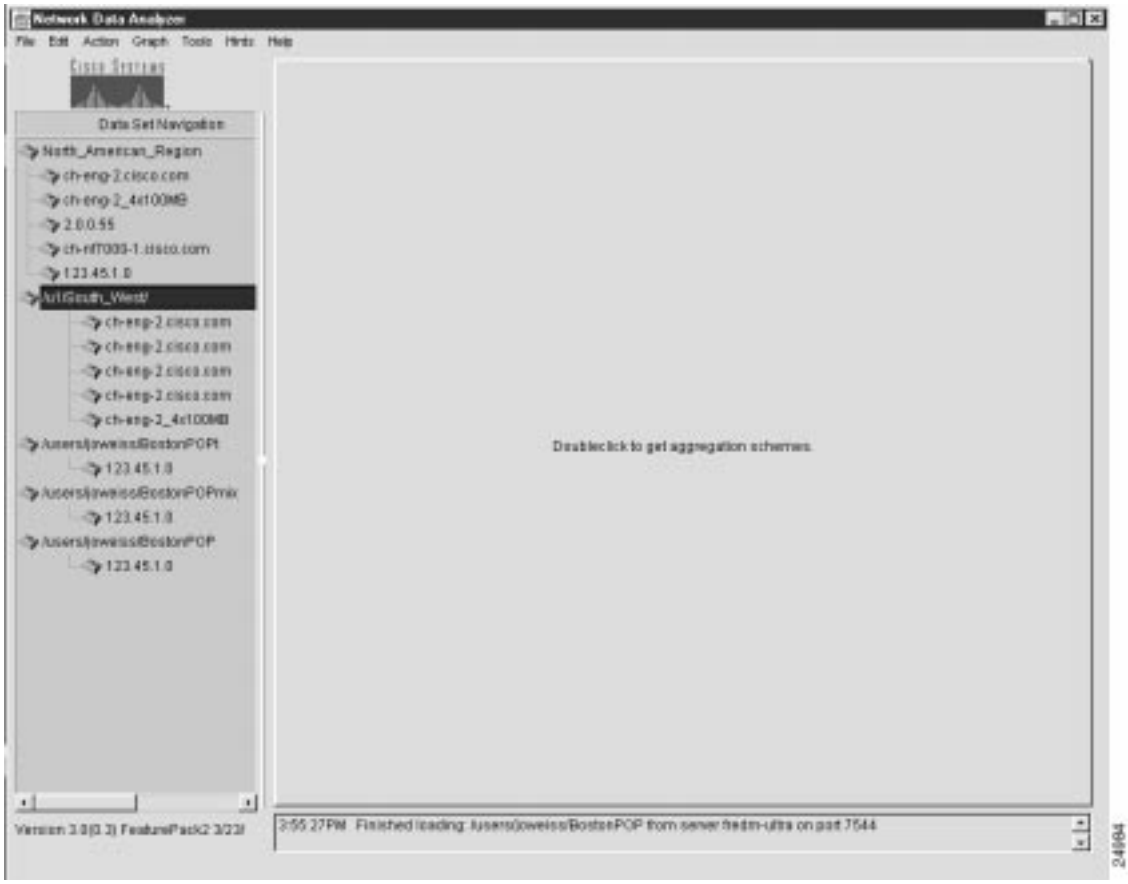

Step 3 Select the Remover Router Or Group option of the Edit menu.

The following pop-up window prompts you for confirmation of your intent to delete the selected router group.

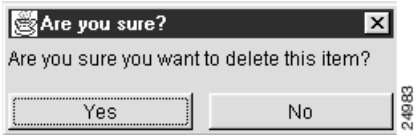

Step 4 Click “Yes.” The Data Set Navigation pane then takes on the appearance shown in Figure 3-14.

Figure 3-14 Effect of Removing a Named Router Group

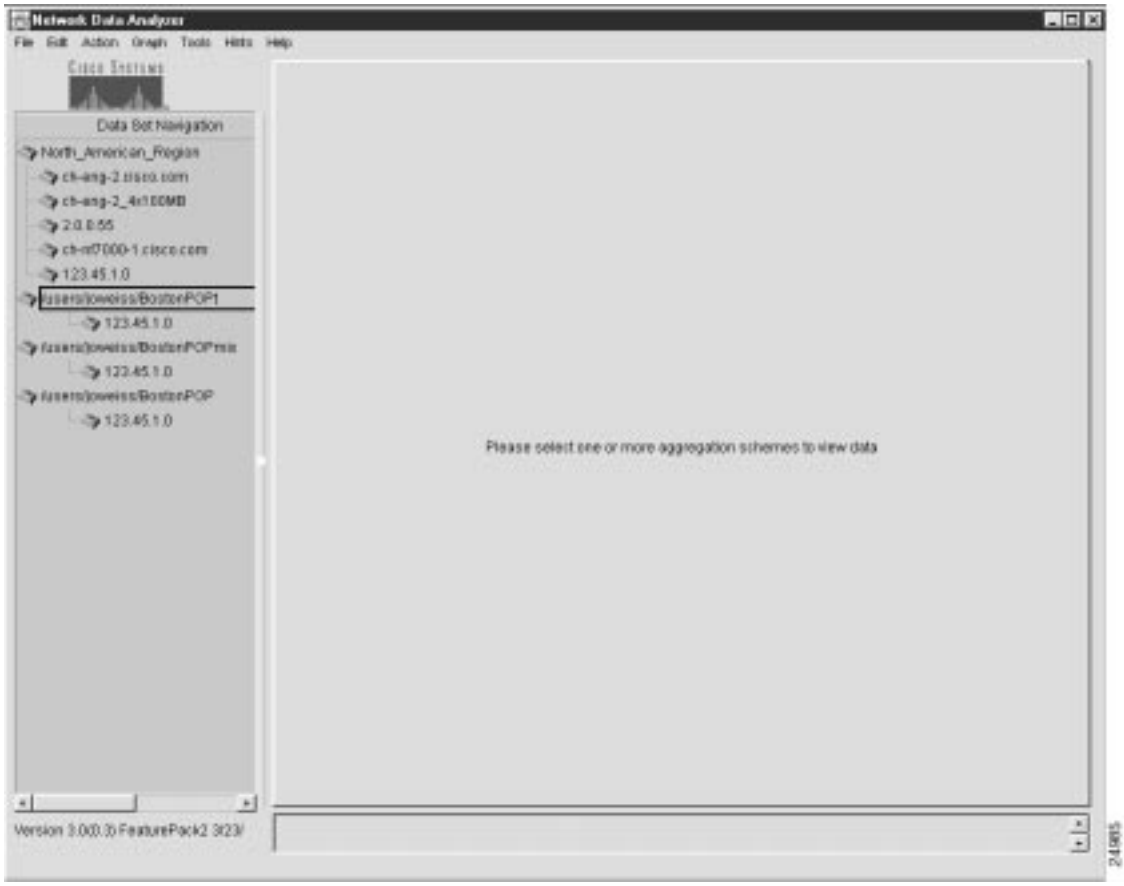

Using the procedure outlined above, you can remove any logical element of an existing data tree structure.

If you select any part of an existing data tree structure for deletion other than one at “root” level (left-justified in the pane), the Remove Router Or Group option remains grayed out. In other words, the Remove Router Or Group option of the Edit menu is effective only with respect to an entire selected router group or selected data set path.

Displaying the Properties of Data Set Paths

By selecting the Properties option of the Edit menu, you can display the properties of any router that you select in the Data Set Navigation pane.

For example, assume that you clicked the router ch-eng-2_4x100MB shown in Figure 3-15 to select it. To display the properties of the selected router, you need only select the Properties option of the Edit menu.

Figure 3-15 Selecting a Data Set Path to Display Its Properties

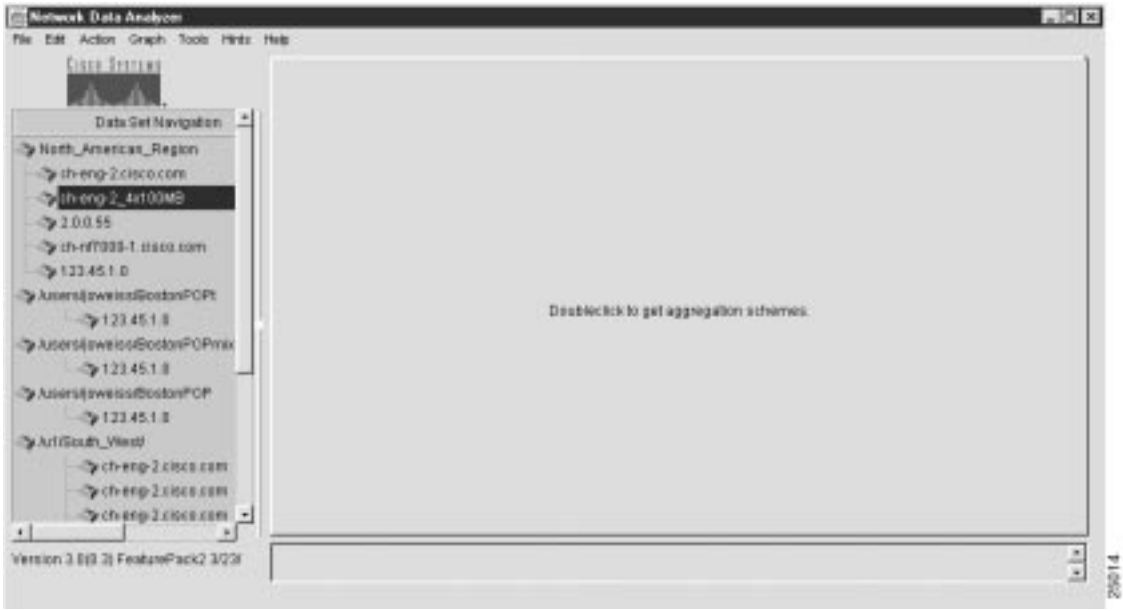

When you select the Properties option of the Edit menu, the following Properties window appears, displaying the properties of the selected router.

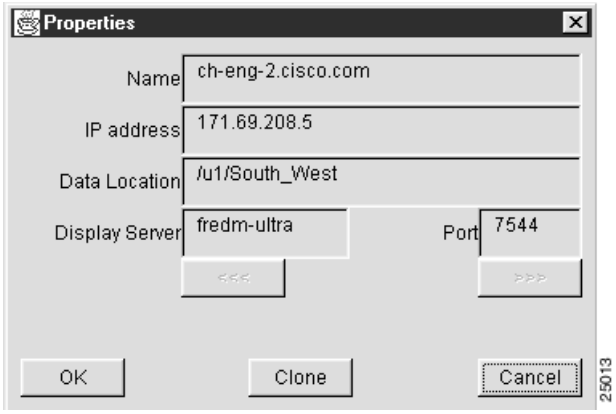

This window displays the following information about the selected router:

- Name of the router
- IP address of the FlowCollector on which NetFlow data for the selected router is being stored
- Name of the FlowCollector directory where the NetFlow data is being stored
- Name of the DisplayServer servicing the Display module
- DisplayServer application port number

The buttons at the bottom of the Properties window perform the following functions:

- OK button—When you click OK, the Properties window is closed. The main Display module window appears.
- Clone button—When you click Clone, a window appears that enables you to clone a router and its associated properties into a selected named router group.

  The procedure for cloning the properties of a selected router into a named router group is described in the following section.
- Cancel—When you click Cancel, the Properties window is closed. The main Display module window appears.

Cloning Function of the Properties Window

The properties window enables you to clone the name and the properties of a selected router into another selected router group.

For example, the cloning process includes the following generalized steps:

1 Select the router “ch-eng-2_4x100MB,” as shown in Figure 3-16, with the intent to clone it into another named router group.

Figure 3-16 Selecting a Router for Cloning Function

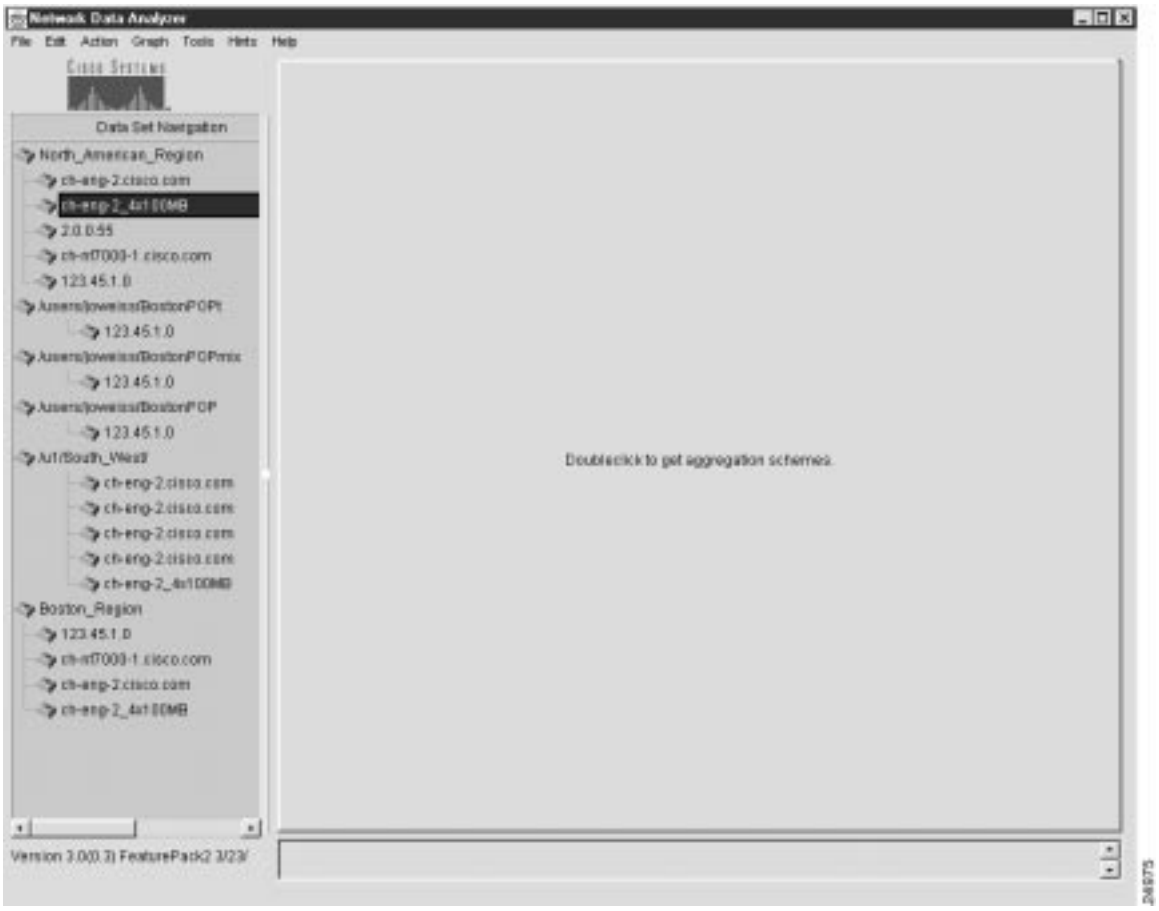

2 With this router selected, select the Properties option of the Edit menu.

This action pops up a Properties window on the main Display module window (see Figure 3-17.

Note You cannot edit the fields of the Properties window. The window displays the properties of the selected router.

Figure 3-17 Invoking the Properties Window

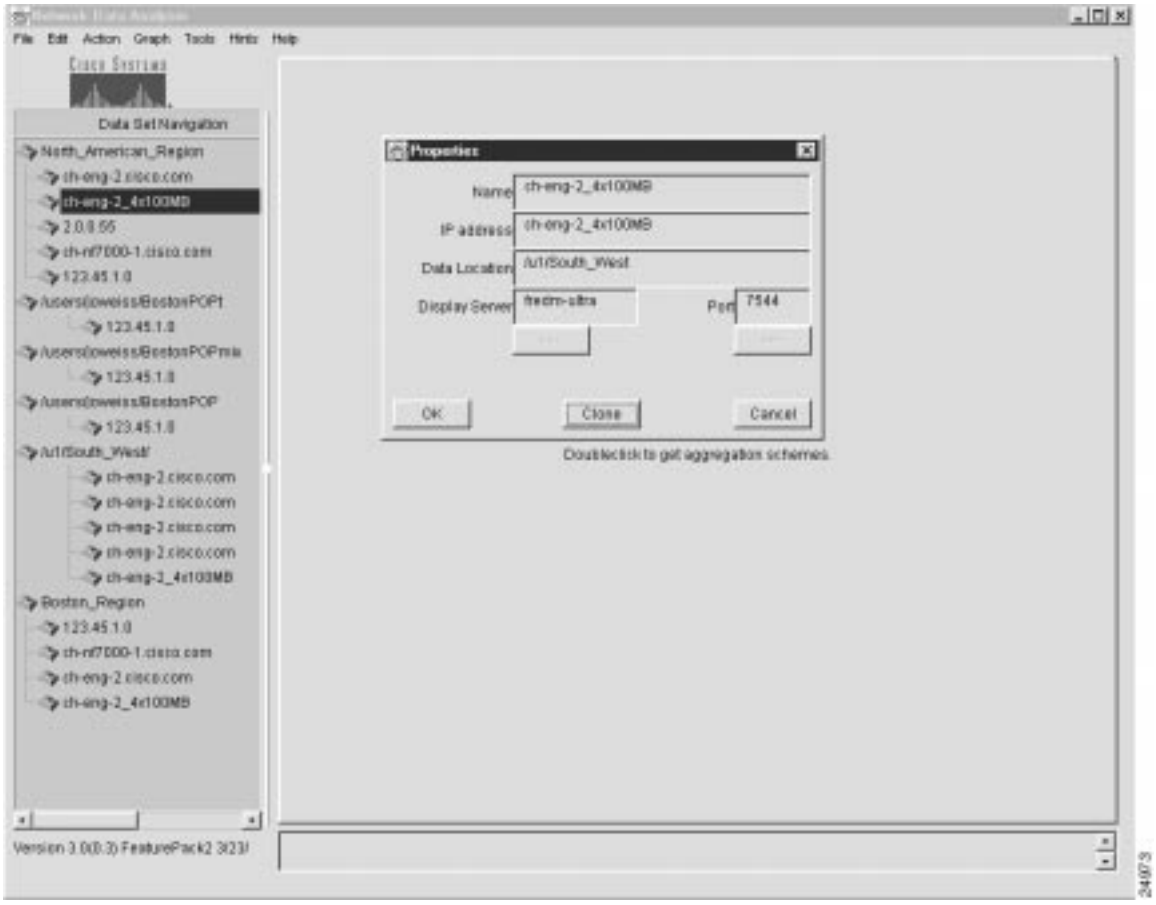

3 Click the Clone button in the Properties window.

This action pops up a Pick a Target Group window on the Properties window (see Figure 3-18).

Figure 3-18 Selecting a Target Router Group for Cloning Operation

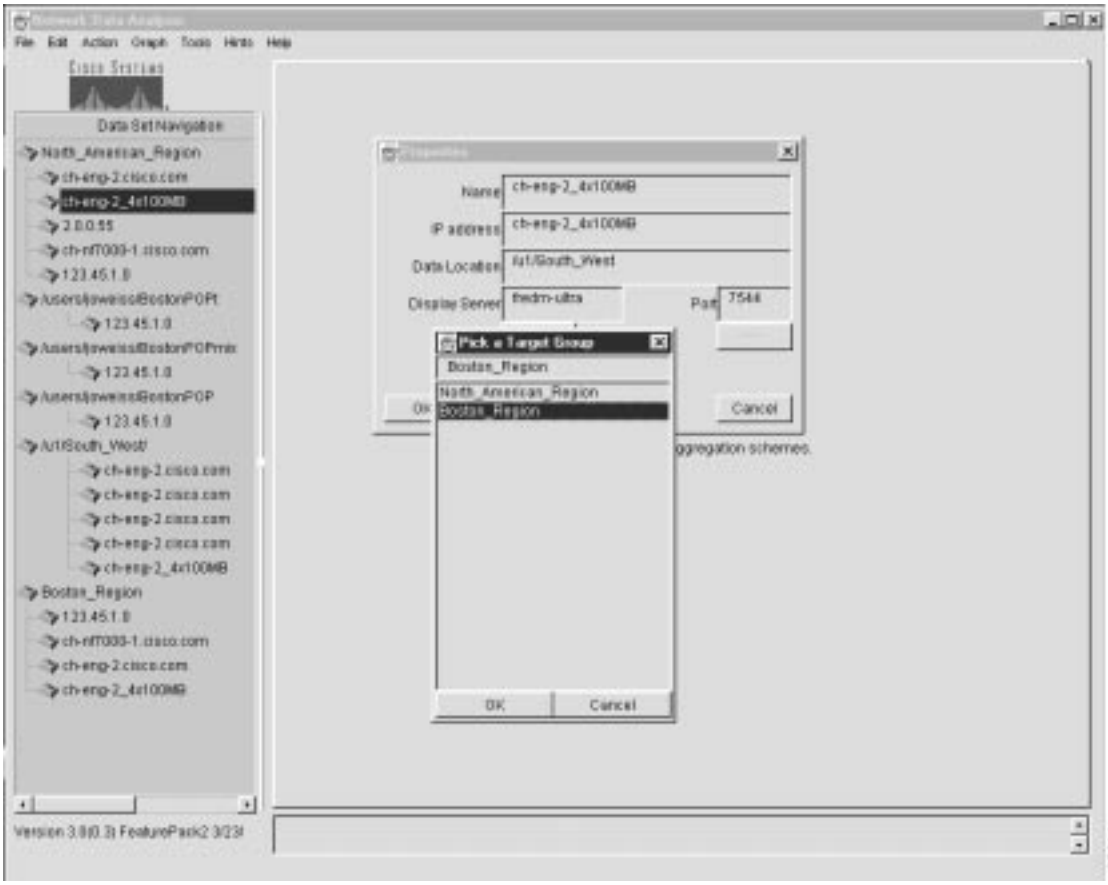

4 In the Pick a Target Group window, select the name of the named router group into which the selected router is to be cloned, such as “Boston_Region.”

5 To execute the cloning function, click Boston_Region (to replicate its name in the field at the top of the window).

6 Click OK.

These steps cause the router named ch-eng-2_4x100MB to be cloned into the router group named Boston_Region, as demonstrated by the new entry at the bottom of the Data Set Navigation pane shown in Figure 3-19.

Figure 3-19 Result of Cloning Operation

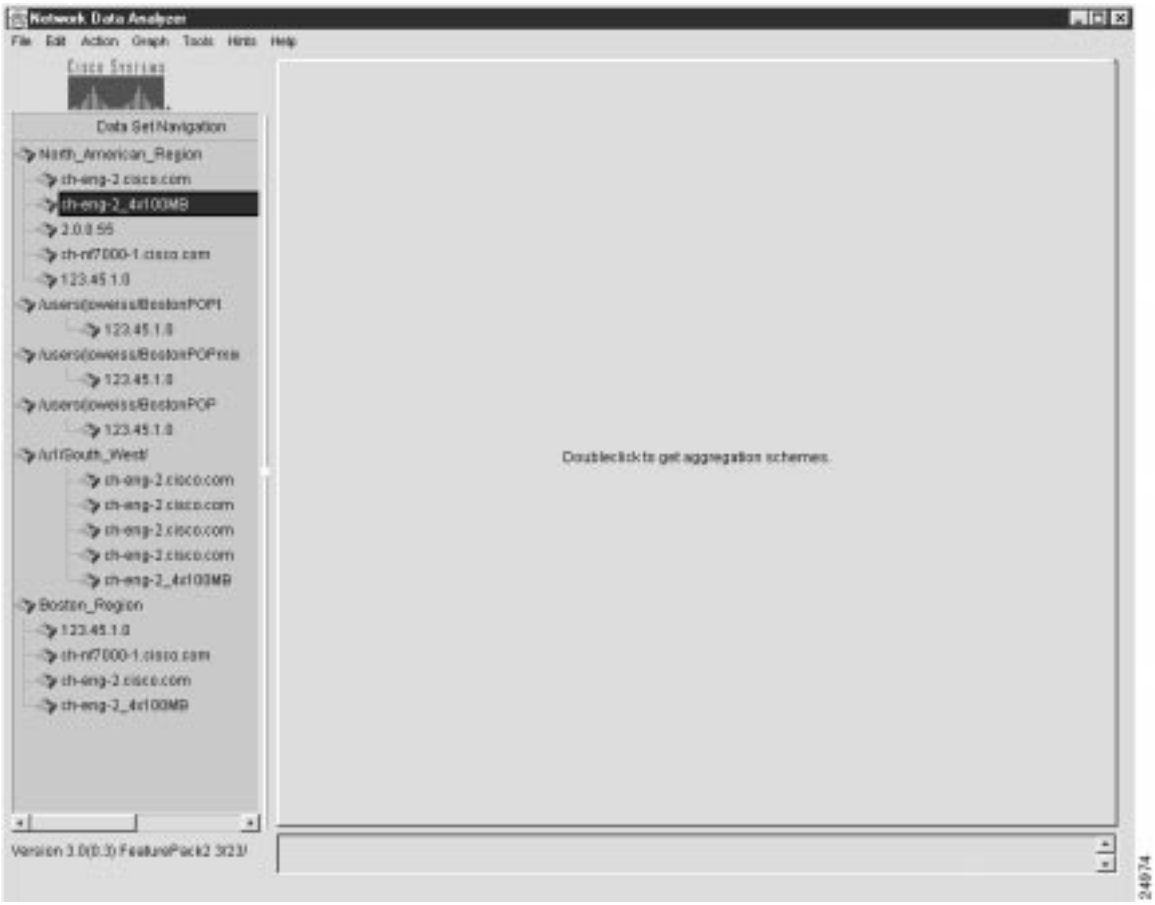

Action Menu Options

The Action menu provides the following selectable options:

- Get Data—See the “Retrieving and Displaying Traffic Data” section on page 3-33.
- Sort Data—See the “Sorting Data for a Selected Aggregation Scheme” section on page 3-39.
- Translate Host Addresses—See the “Translating Host IP Addresses” section on page 3-42.

The following sections describe how to use these menu options.

Retrieving and Displaying Traffic Data

This section presents procedures for retrieving and displaying traffic information for:

- A selected aggregation scheme
- A selected device
- A named router group that has a common aggregation scheme

Displaying Data for a Selected Device

A common Analyzer task is to display traffic information for a selected device and aggregation scheme.

For purposes of this section, assume that you have added the data tree structure shown in Figure 3-20 to the Data Set Navigation pane.

Figure 3-20 Sample Tree Structure for Displaying Selected Aggregation Scheme

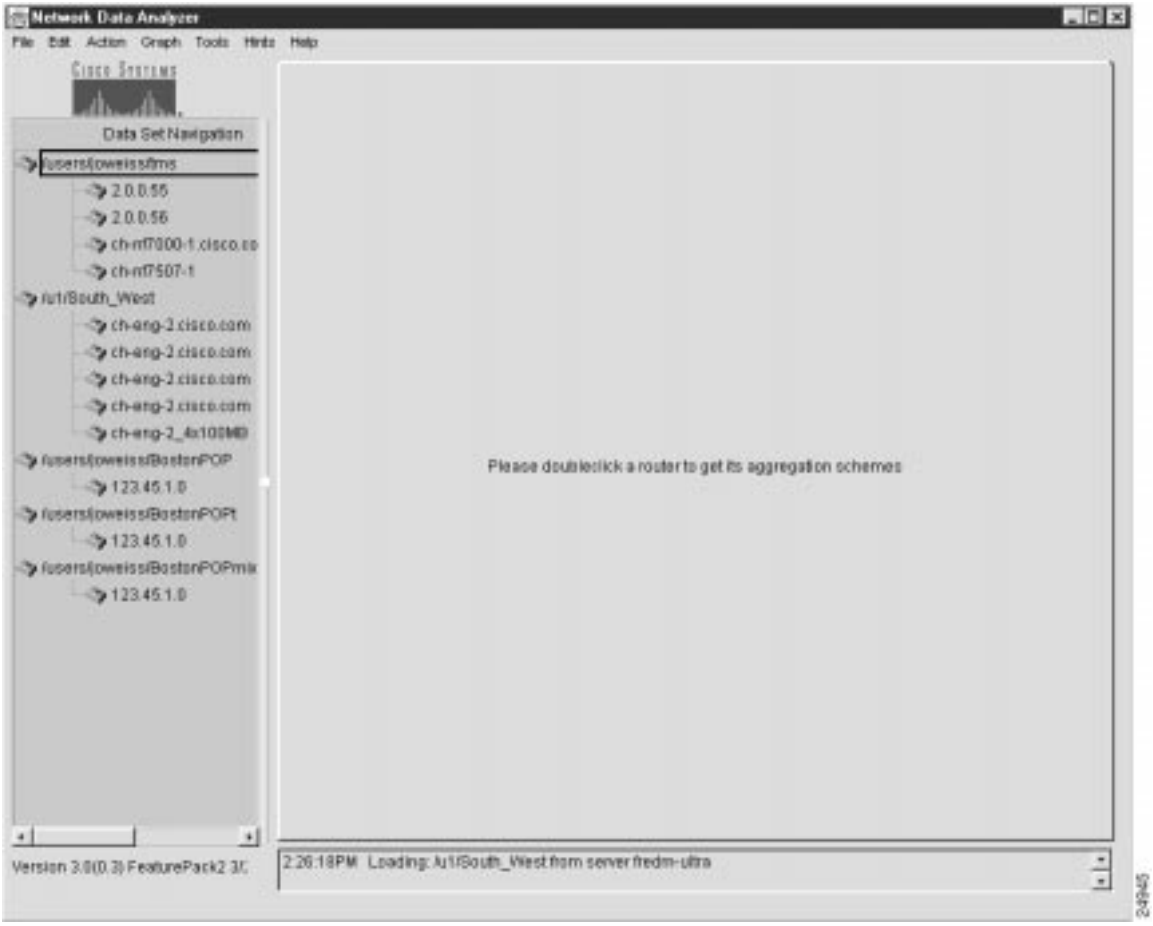

To display traffic data for a selected device and aggregation scheme, perform the following steps:

Step 1 Select the desired device.

Double-click the name of the router of interest in the Data Set Navigation pane (see Figure 3-20).

Assuming that “ch-eng-2_4x100MB” is the device of interest, double-click the device name to list its applicable aggregation schemes in the Data Set Navigation pane.

Step 2 Select the aggregation scheme that you want to apply for display purposes.

For example, assume that you want to display traffic information for the DetailInterface aggregation scheme. Click this aggregation scheme.

This action causes a “clean” display pane labeled DetailInterface to appear in the Display module window (see Figure 3-21).

Figure 3-21 Display Pane for DetailInterface Aggregation Scheme

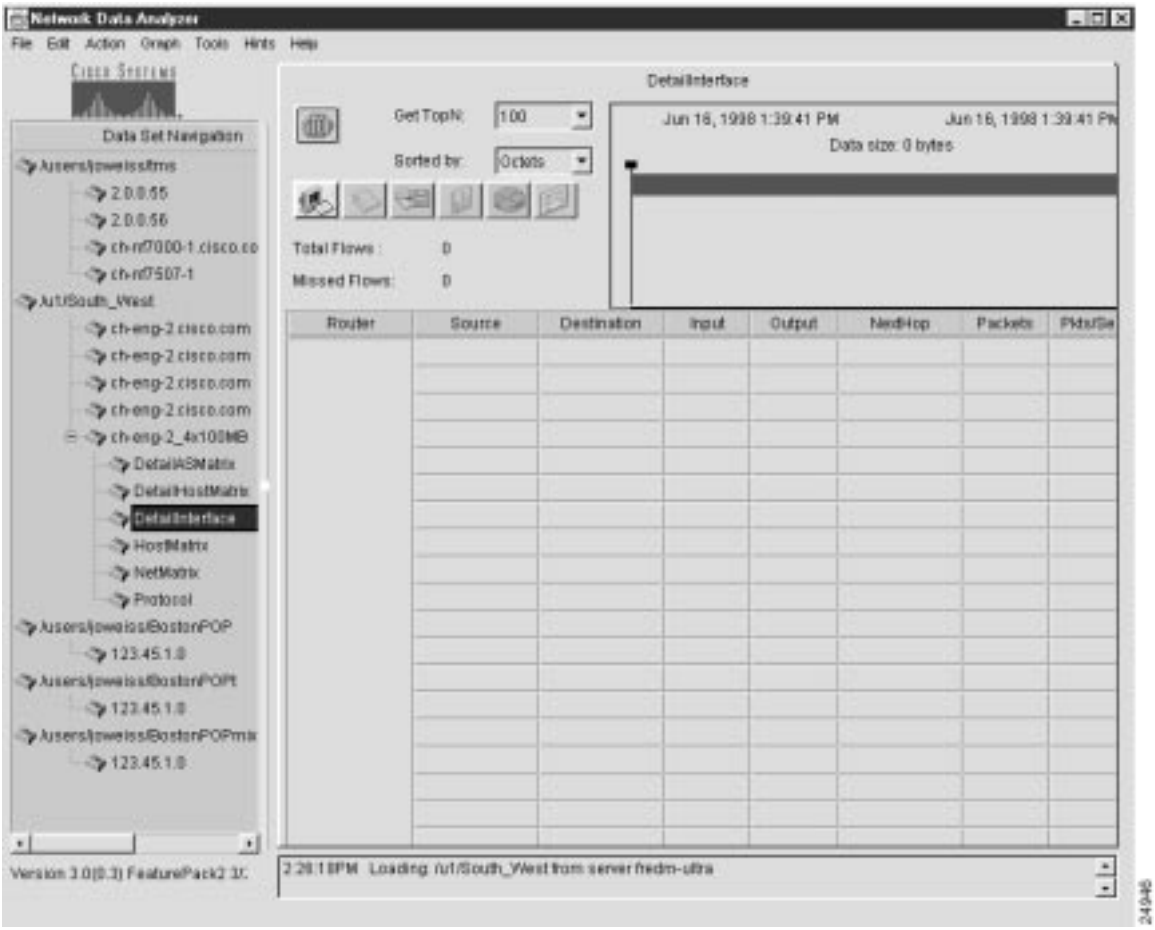

Note that the time line area in the top right portion of the window indicates the range of dates and times for which data is available for the selected router and aggregation scheme.

Step 3 To establish the desired time horizon for data retrieval purposes, position the time slider marks, as appropriate.

You can move either time slider mark in either horizontal direction to establish the desired time period.

Step 4 Click the Get TopN: pull-down menu to select the number of traffic flows that you want taken into account in data retrieval operations.

The selectable values of "N" range from 10 to 10,000. The value 100 is the default.

In effect, the value that you select determines the relative volume of traffic data to be retrieved and processed for the selected aggregation scheme.

For example, if you select "10" as the TopN value, you limit the volume of traffic data to be processed to the first ten traffic flows.

Selecting a lower value for N tends to improve Analyzer performance, because less information is processed in satisfying the data display request.

Step 5 Click the Sorted by: pull-down menu to select the desired sort key for displaying the traffic data.

For this purpose, assume that the data is to be sorted by “Octets.”

Step 6 Click the Get Data from Server button in the display pane (the leftmost button).

Alternatively, you can select the Get Data option from the Action menu to initiate the retrieval and display of traffic data for the selected aggregation scheme.

The result of this procedure is shown in Figure 3-22.

Figure 3-22 DetailInterface Aggregation Scheme for the Selected Device

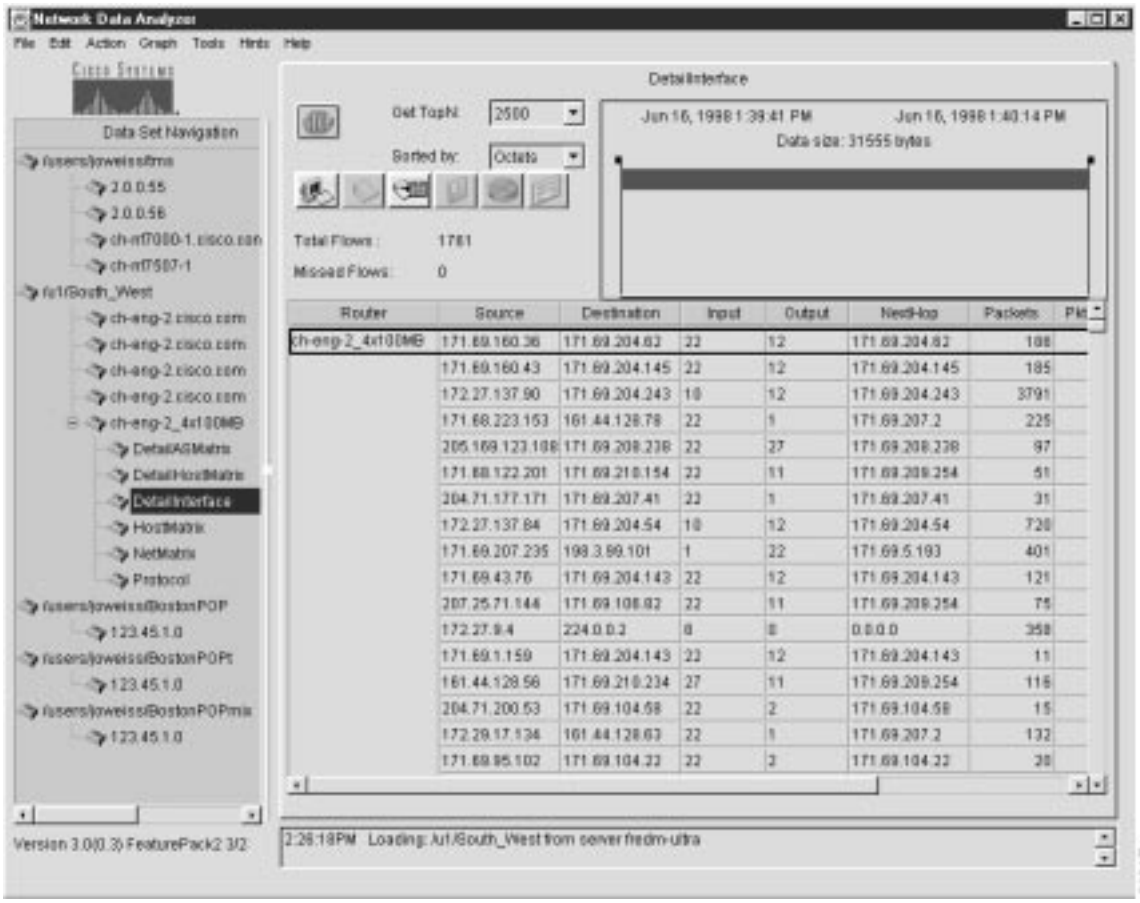

Displaying Data for Aggregation Scheme Common to Devices in Named Router Group

A common Analyzer task is to display traffic information for a selected aggregation scheme that is common across devices in a named router group.

Figure 3-23 shows the tree structure for the named router group, “/u1/South_West,” that serves as the basis for the procedures described in this section. Figure 3-23 indicates that you want to display NetFlow data for the DetailASMatrix aggregation scheme for three exporting devices that are common to the named router group.

For details about how to create a named router group, refer to the “Creating and Saving a Named Router Group” section on page 3-20.

Figure 3-23 Tree Structure for a Named Router Group

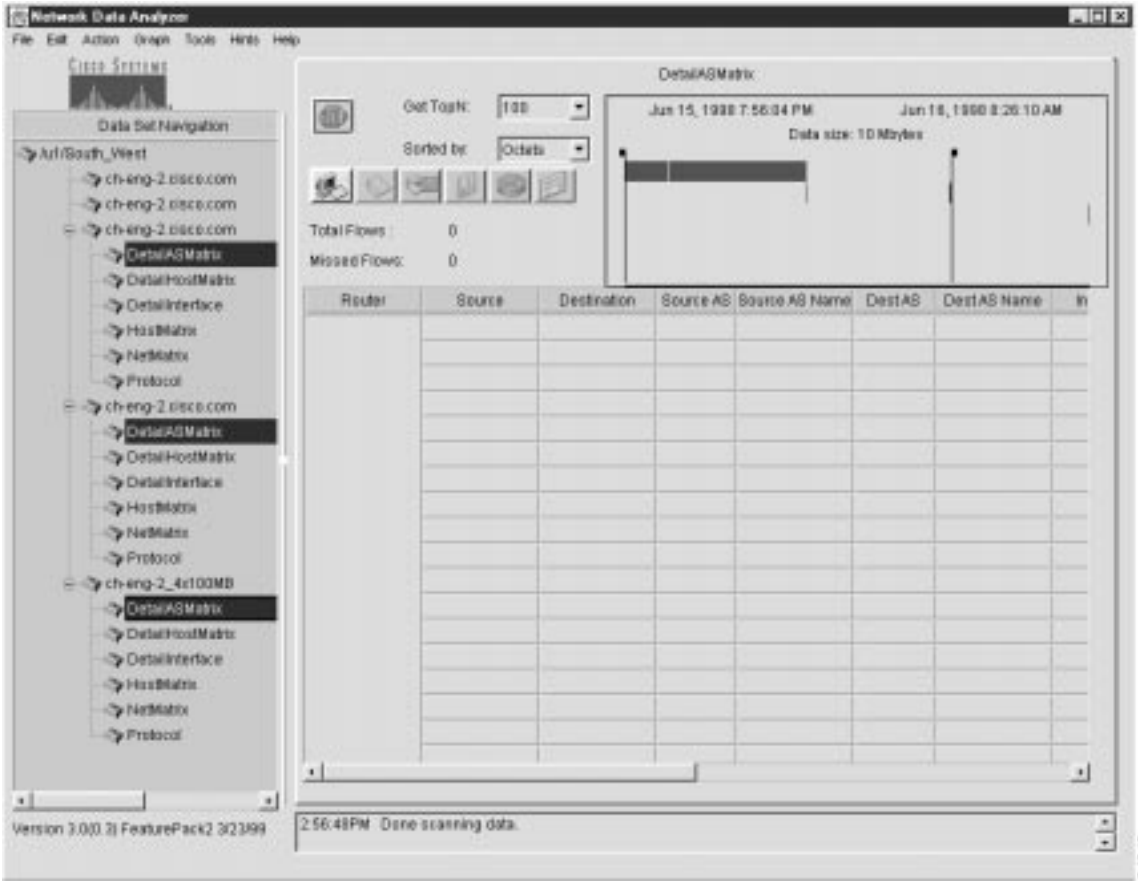

To display NetFlow data for an aggregation scheme that is common to selected devices in a named router group, perform the following steps:

Step 1 Select a desired aggregation scheme that is common to two or more devices in the overall tree structure for the named router group. The selected aggregation scheme must be the same for all devices.

To select a common aggregation scheme among two or more devices, hold down the **Ctrl** key and click the name of each scheme.

Step 2 Position the time slider marks in the display pane, as desired, to establish the applicable time horizon for display operations.

Step 3 Click the Get TopN: pull-down menu in the display pane to select the number of traffic flows (“N”) that you want taken into account for display purposes. The default value for Get TopN: is 100.

Step 4 Click the Sorted by: pull-down menu in the display pane to select the desired sort key for display purposes. The default value for the Sorted by: parameter is “Octets.”

Step 5 Click the Get Data from Server button in the display pane (the leftmost button).

Alternatively, you can select the Get Data option from the Action menu to initiate the display operation for the selected aggregation scheme.

Note If the volume of data for the selected aggregation scheme exceeds the memory capacity of your Display module host, the system may report that it has run out of memory. In this case, restart the Display module. If you again attempt to display aggregation scheme data for multiple network devices, you can do one or all of the following, depending on the availability of memory: a) reduce the number of devices to be included in the display operation; b) specify a smaller value for N (the number of flows); or c) reduce the applicable time horizon for the display operation.

On completion of the above procedure, traffic data for the selected aggregation scheme appears (see Figure 3-24).

Figure 3-24 Aggregation Scheme Data for Selected Devices in Named Router Group

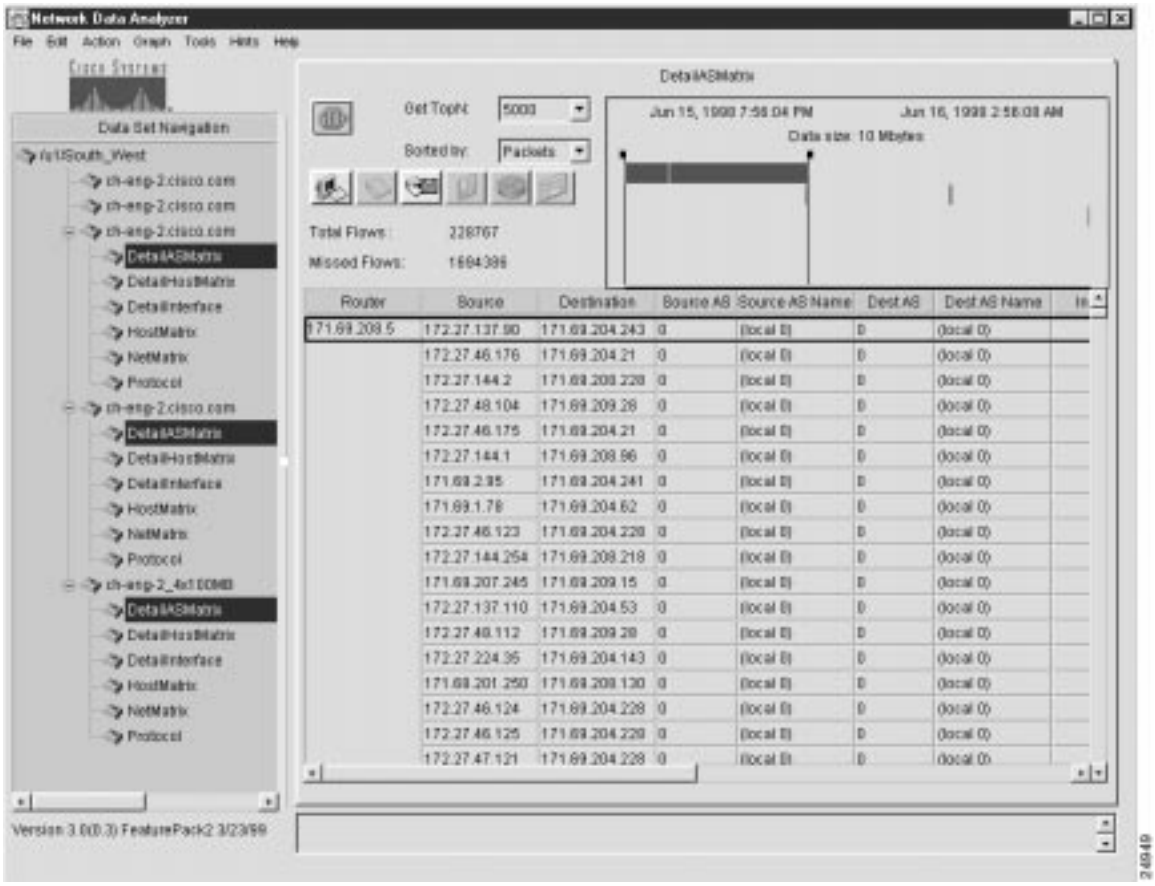

The applicable aggregation scheme data for each selected device is “stacked” in the display pane one below another, beginning with the first selected device and continuing to the last selected device.

As with all display functions, when the overall data array exceeds the observable “at-a-glance” space in the display pane, you can use the horizontal and vertical scroll bars in the display pane to traverse to any area of the display.

By using similar data tree structures in this manner for multiple devices in a named router group, you can select any aggregation scheme common among any combination of devices in the group and initiate display functions for the selected devices. This capability enables you to compare traffic patterns and statistics of interest between exporting devices in the group.

Sorting Data for a Selected Aggregation Scheme

For purposes of this section, the working data tree structure shown in Figure 3-25 is assumed to be basis for sort operations.

Figure 3-25 Sample Data Tree Structure for Sort Operations

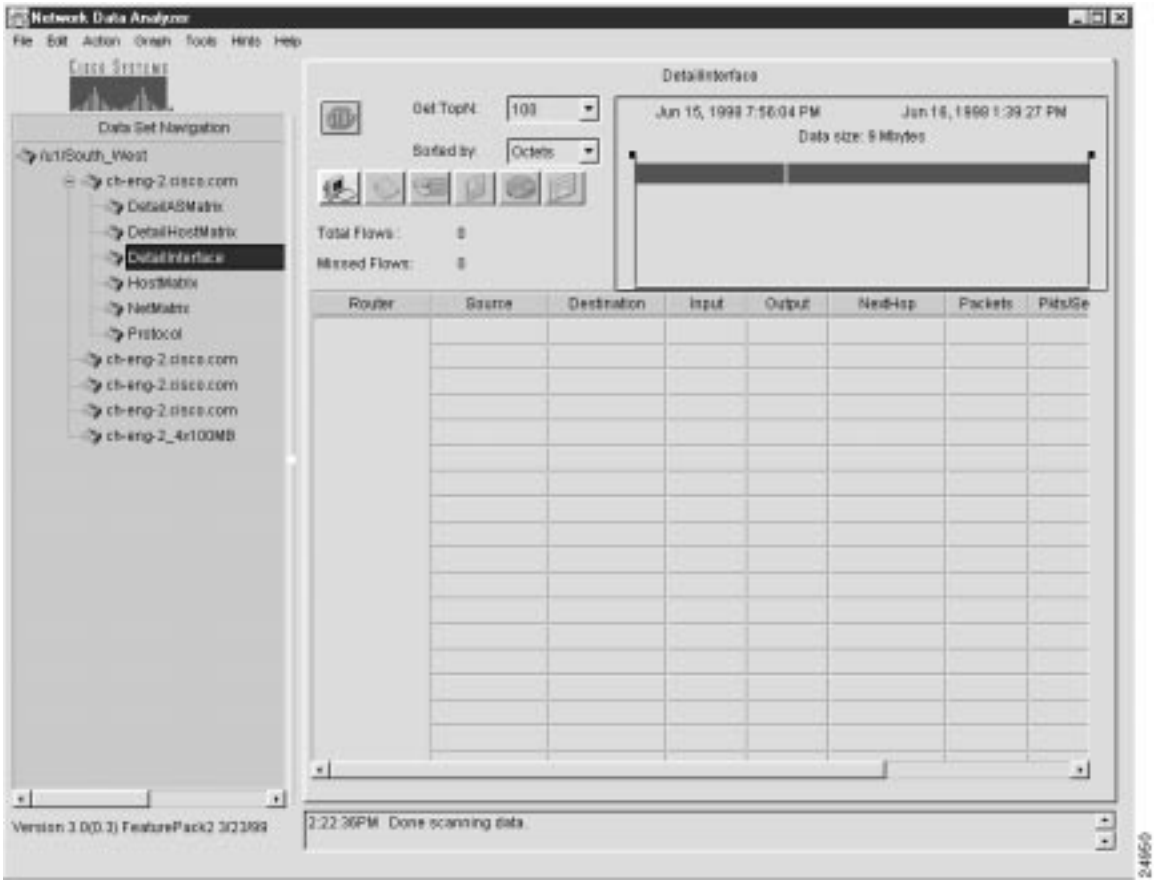

You can select any aggregation scheme in an existing data tree structure, retrieve the applicable traffic data for that aggregation scheme, select any column of the data array, and sort the data in that column.

To initiate a sort operation relative to a selected aggregation scheme, perform the following steps:

Step 1 Select an aggregation scheme for display purposes.

For this step, assume that you want to display traffic information for the DetailInterface aggregation scheme.

Step 2 Click the DetailInterface aggregation scheme to access a "clean" display pane.

Step 3 Prepare for sort operations, as follows:

(a) Position the time slider marks in the display pane as desired to establish the applicable time horizon for governing the extent of data retrieval.

(b) Click the Get TopN: pull-down menu in the display pane to select the number of traffic flows ("N") that you want taken into account for data retrieval purposes. The default value for Get TopN: is 100.

(c) Click the Sorted by: pull-down menu to select the desired sort key for display purposes. The default value for the Sorted by: parameter is "Octets."

Step 4 Click the Get Data from Server button in the display pane.

Alternatively, you can select the Get Data option from the Action menu to initiate data retrieval and display for the selected aggregation scheme.

Step 5 When the data for the selected aggregation scheme appears in the display pane (Figure 3-26), you can sort the data in either of two ways:

- Option 1—Click the heading of any column of interest in the data array to highlight that column, then:
  - Select the Sort Data option from the Action menu to initiate the sort operation, or
  - Click the Sort data selected button in the display pane window to initiate the sort operation.
- Option 2—Double-click any column heading of interest in the data array. This will initiate the sort operation for that column.

  Using this option, you can double-click a succession of column headings (in any desired order) to initiate sort operations for a series of selected columns.

For purposes of this step, assume that you want to sort the data in the Packets column of the data array.

Figure 3-26 Sample Aggregation Scheme Data for Sort Operations

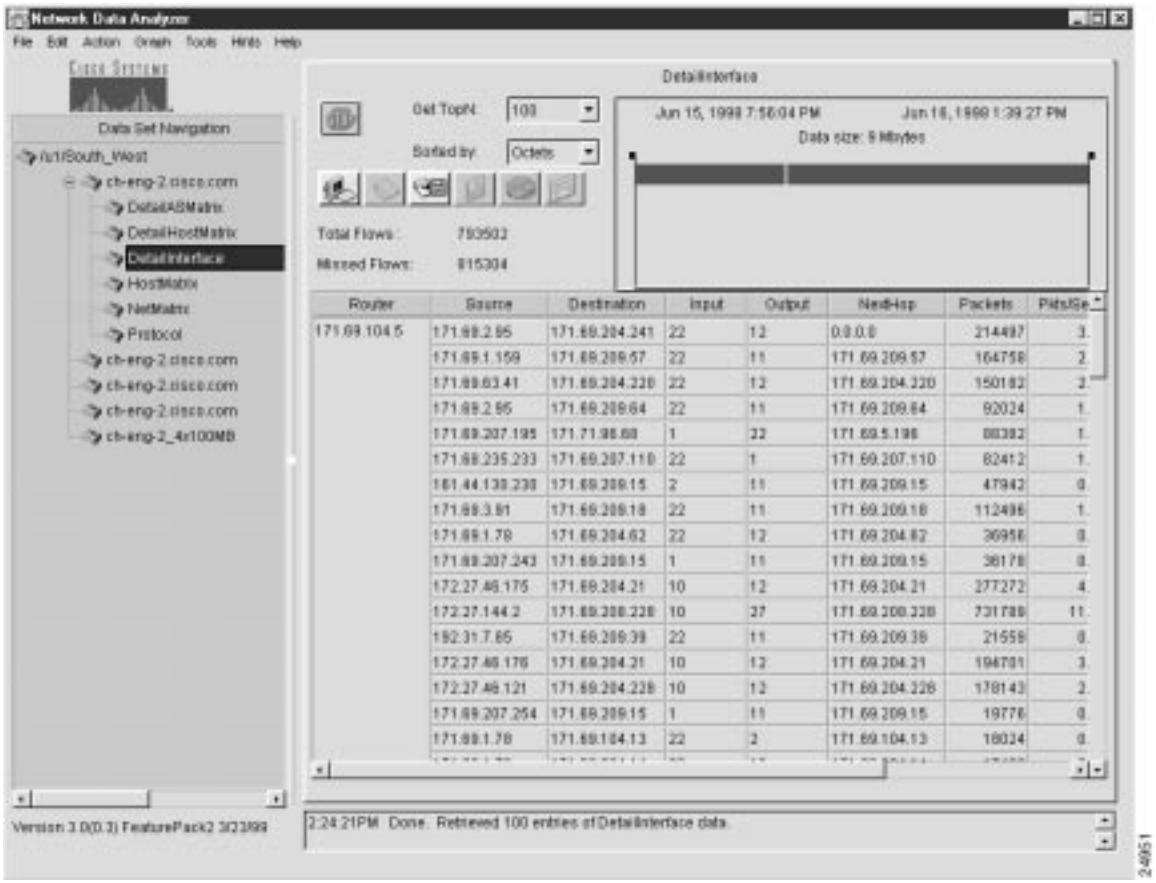

This procedure results in the data array shown in Figure 3-27.

Figure 3-27 Sort Operations on Packets Column of Selected Aggregation Scheme

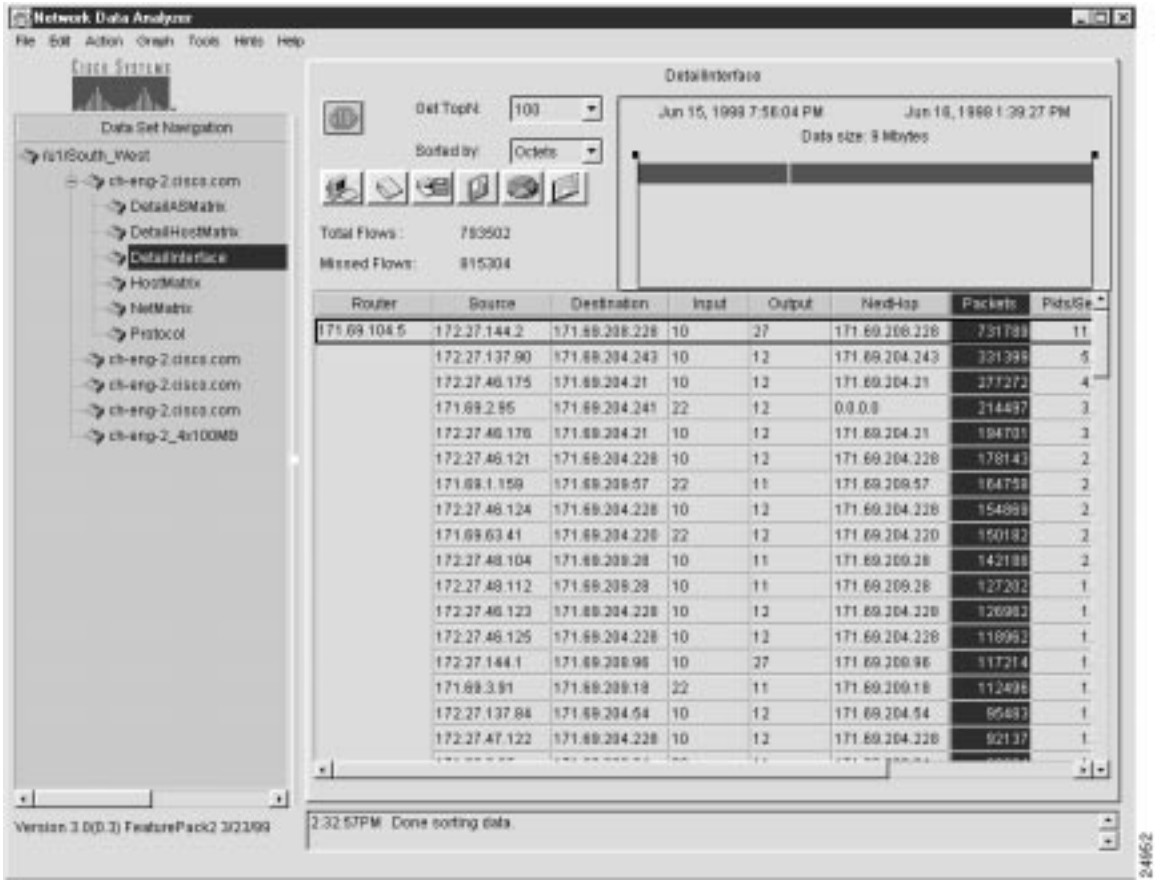

Note that the data in the Packets column of the aggregation scheme is sorted in descending numerical order.

Translating Host IP Addresses

The Translate Host Addresses option of the Action menu enables you to translate IP addresses appearing in the Source or Destination columns of a displayed NetFlow data aggregation scheme into equivalent host names.

Table 3-1 lists the NetFlow data aggregation schemes and the key columns applicable to each scheme. For purposes of translating host IP addresses, the applicable key columns of a NetFlow data aggregation scheme include only the following:

- srcaddr—Source IP address (the first key column in Table 3-1)
- dstaddr—Destination IP address (the second key column in Table 3-1)

The Translate Host Addresses option is applicable only to NetFlow data aggregation schemes that incorporate Source or Destination columns in the displayed aggregation scheme. If the Translate Host Addresses option is not available, it is grayed out.

When you invoke the Translate Host Addresses option, the NetFlow data array in the display pane changes, as follows:

- For the Source column—A new column labeled Source Name is incorporated into the data array immediately to the right of the Source column. It contains the host names of the associated source devices.
- For the Destination column—A new column labeled Dest Name is incorporated into the data array to the right of the Destination column. It contains the host names of the associated destination devices.

Note that the following Display module functions incorporate an IP address translation function:

- AS Drill Down Window option of the Tools menu—This window incorporates a Translate button that is operationally equivalent to the Translate Host Addresses option of the Action menu. For more information, refer to the "Drilling Down on Network Flows" section on page 3-50.
- Search Window of the Tools menu—This window incorporates a Translate button that is also operationally equivalent to the Translate Host Addresses option of the Action menu. For more information, refer to the "Searching for Flows by Source and Destination Addresses" section on page 3-57.

Figure 3-28 shows a typical NetFlow data aggregation scheme containing Source and Destination columns.

Figure 3-28 Sample Aggregation Scheme for Translate Host Addresses Function

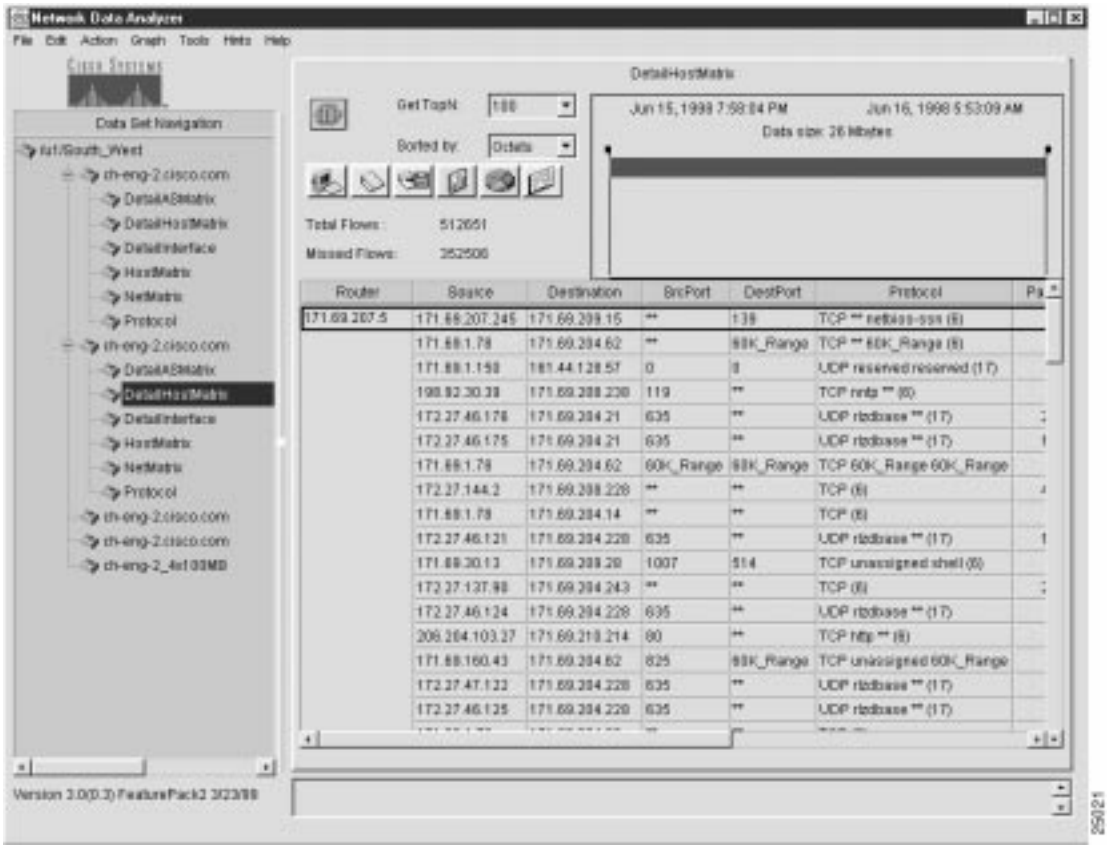

If you invoke the Translate Host Addresses option while this aggregation scheme is in effect, the data array is altered as follows:

- Two new columns would be added to the data array:
  - Source Name—This new column appears to the right of the existing Source column.
  - Dest Name—This new column appears to the right of the existing Destination column.
- The equivalent Source and Destination host names are translated and posted to the newly-created Source Name and Dest Name columns of the altered display pane.

The result of this IP translation function is shown in Figure 3-29.

Figure 3-29 Result of Translating IP Addresses into Host Names

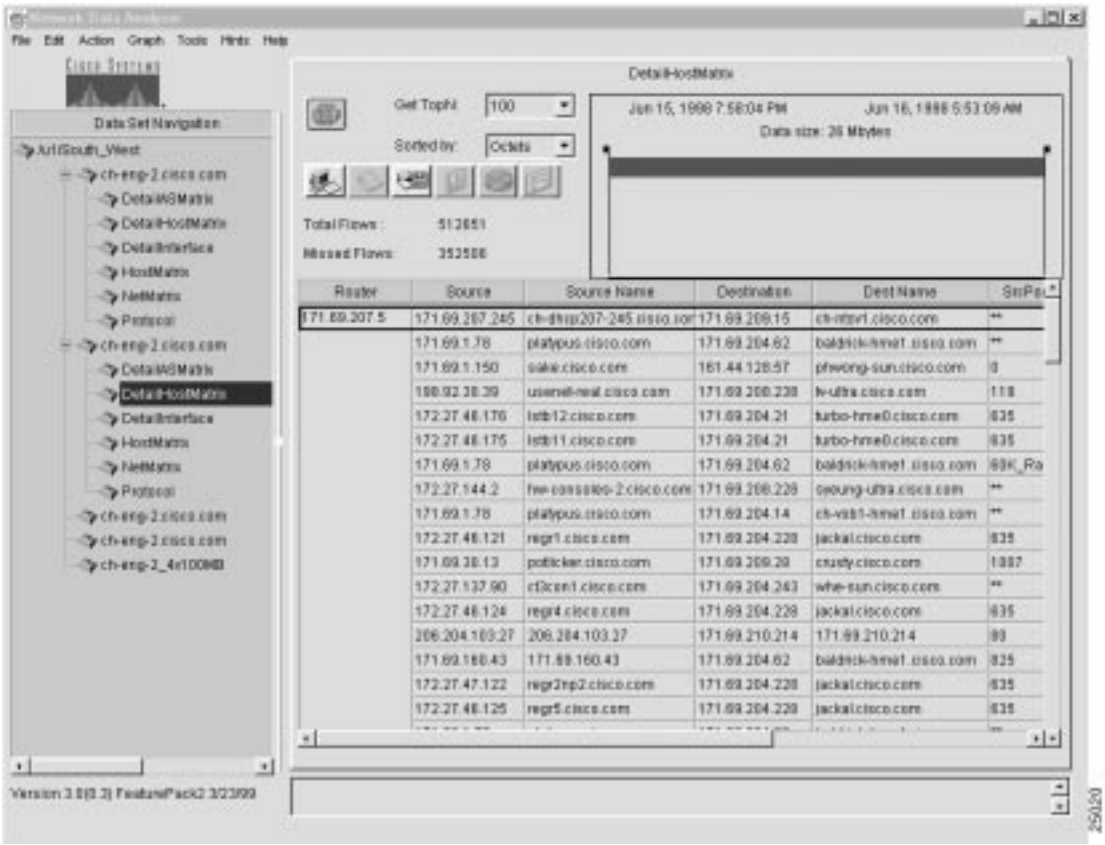

You may find it useful to sort the host names appearing in the Source Name and Dest Name columns of a NetFlow data array.

For example, if several exporting devices appear in your network that are identified with host names that vary slightly from each other (such as fredm-ultra, fredm-sun, fredm-pc), you can perform a sort on the Source Name or the Dest Name column in the data array. In this example, the sort operation causes the “fredm” host devices to sort the data alphabetically and rearrange it in the resulting data array.

Graph Menu Options

The Display module Graph menu provides the following selectable options:

- Bar—See the “Creating a Bar Chart” section on page 3-45.
- Pie—See the “Creating a Pie Chart” section on page 3-47.
- Histogram—See the “Creating a Histogram” section on page 3-48.

The following sections describe how you use these menu options.

Creating a Bar Chart

The data in the DetailInterface aggregation scheme shown in Figure 3-30 is used as the basis for illustrating the Analyzer’s bar chart function.

Figure 3-30 Typical Aggregation Scheme Data for Bar Chart Creation

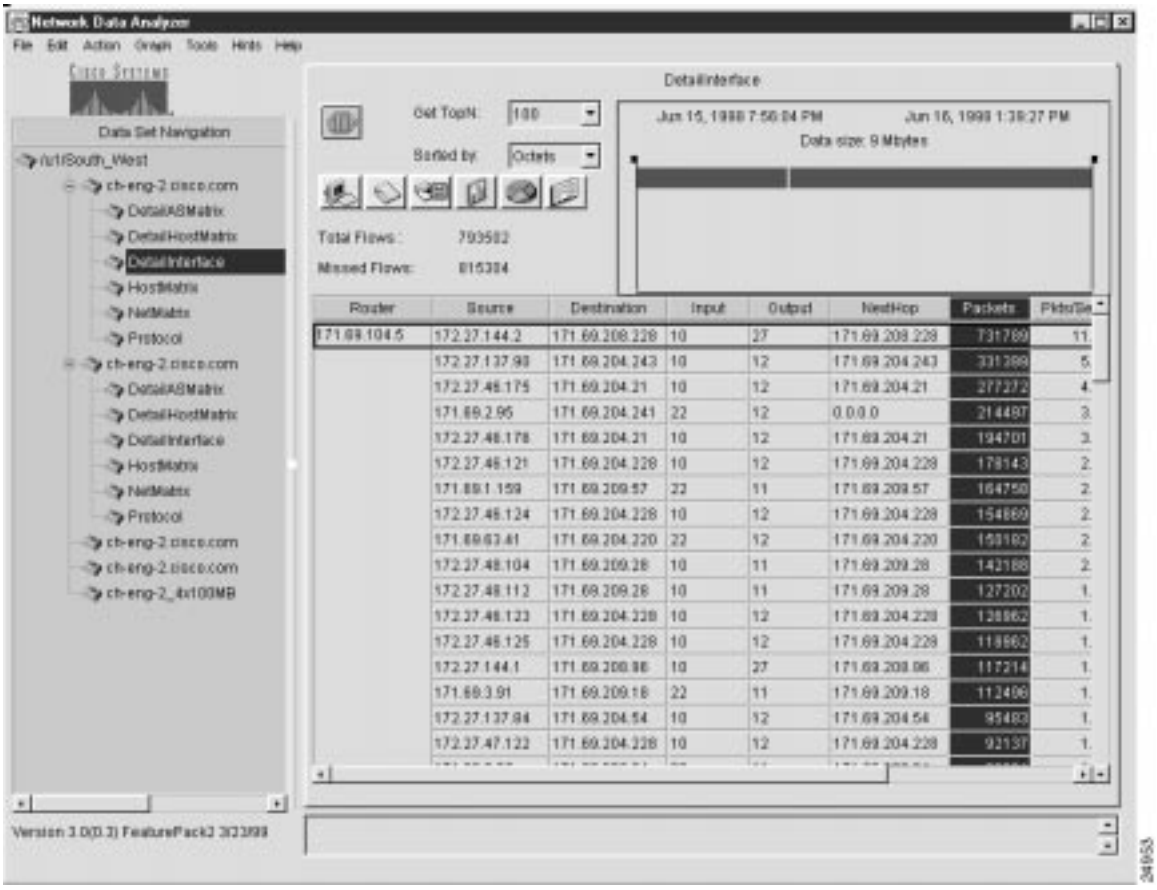

The Packets column of the aggregation scheme is highlighted, indicating that you can use the data in this column for bar chart creation. You can select any of the six value columns of a NetFlow data aggregation scheme, or any of the 14 value columns of a TMS data aggregation scheme as the basis for creating a bar chart.

You can use either of two methods to create a bar chart representation of the traffic information in the Packets column (or any other value column of a NetFlow or TMS data aggregation scheme):

- Method 1—Click the Draw Bar Chart button in the display pane. This action pops up a Graph TopN pop-up window.

  Place the mouse pointer in this pop-up window. A list of choices for the Graph TopN value appears.

  Click the desired value to initiate bar chart creation. For example, assume that “30” is the value of choice.

  Selecting “30” as the Graph TopN value causes the 30 largest traffic flows (in terms of packets) to be incorporated into the horizontal axis of the bar chart.

  The resulting bar chart (see Figure 3-31) appears in a separate window on your Display module screen.

  To remove the bar chart from the screen, click Close.

Figure 3-31 Bar Chart for Selected Column of Aggregation Scheme

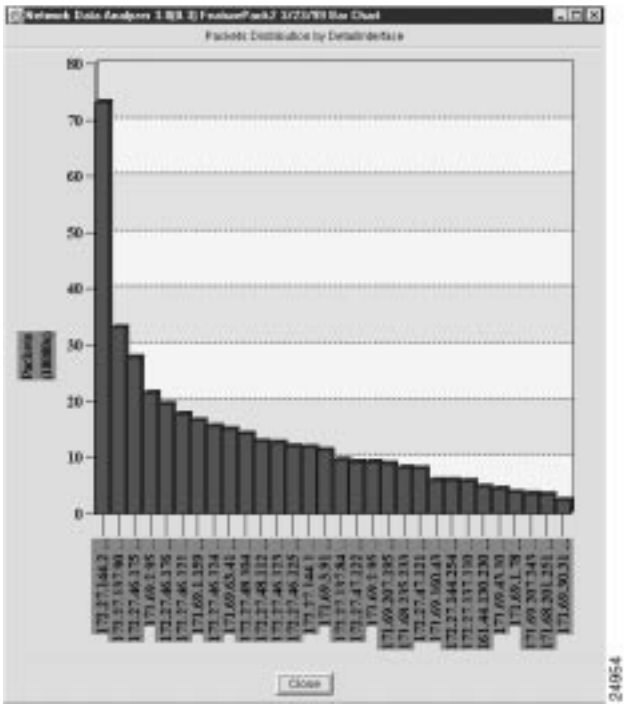

- Method 2—Click the Graph pull-down menu of the Display module window. This action pops up a window that lists Bar, Pie, or Histogram as options.

  Place the mouse pointer on Bar. A list of choices for the TopN value appears. Move the mouse pointer to the desired value (assumed to be “30” for this step) and click the value to initiate bar chart creation.

  The resulting bar chart (Figure 3-31) appears in a separate window on your Display module screen.

  To remove the bar chart from the screen, click Close.

Creating a Pie Chart

For this section, the aggregation scheme data shown in Figure 3-32 serves as the basis for illustrating pie chart creation.

Figure 3-32 Typical Aggregation Scheme Data for Pie Chart Creation

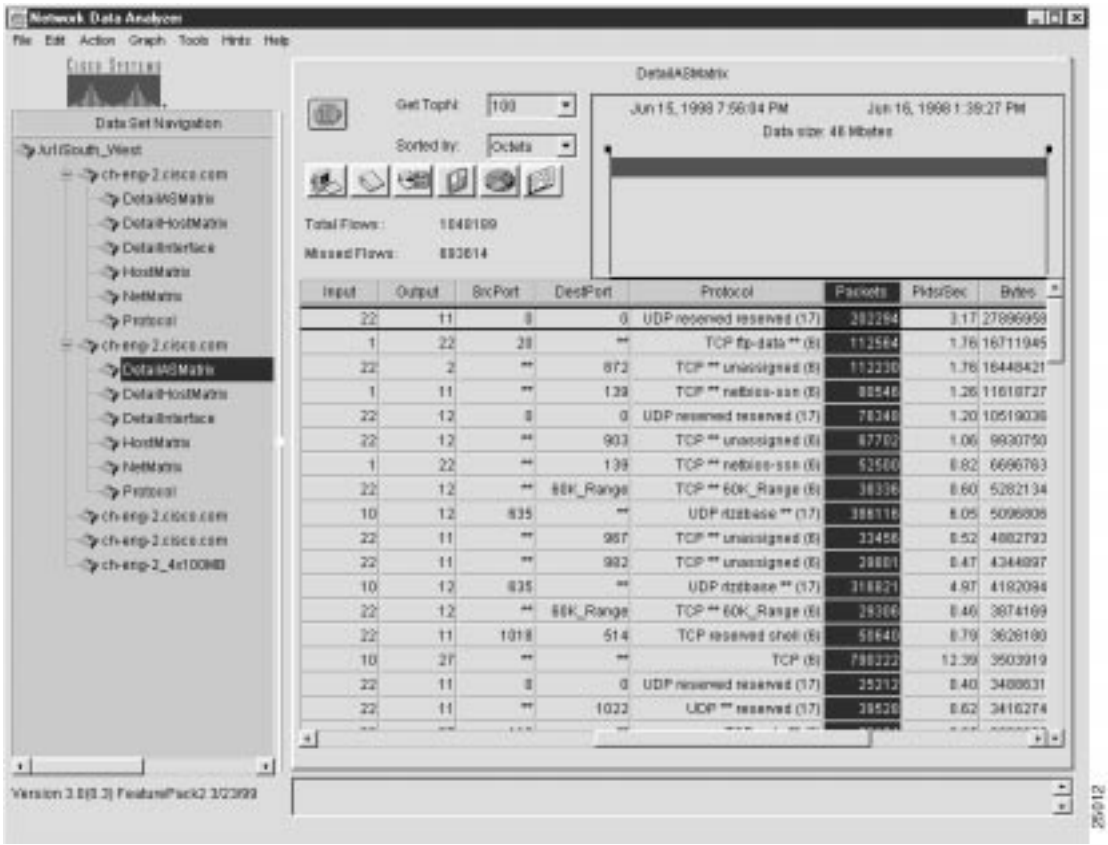

The Packets column of the aggregation scheme is highlighted, indicating that you can use the data in this column for pie chart creation. As with bar charts, you can select any of the six value columns of a NetFlow data aggregation scheme, or any of the 14 value columns of a TMS data aggregation scheme as the basis for creating a pie chart.

You can use either of two methods to create a pie chart representation of the traffic information in the Packets column (or any other value column of a NetFlow or TMS data aggregation scheme):

- Method 1—Click the Draw Pie Chart button in the display pane. This action pops up a Graph TopN pop-up window.

  Place the mouse pointer in this pop-up window. A list of choices for the Graph TopN value appears.

  Click the desired value to initiate pie chart creation. For example, assume that “30” is the value of choice.

  Selecting “30” as the Graph TopN value causes the 30 largest traffic flows (in terms of packets) to be incorporated into the horizontal axis of the pie chart.

The resulting pie chart (see Figure 3-33) appears in a separate window on your Display module screen.

To remove the pie chart from the screen, click Close.

Figure 3-33 Pie Chart for Selected Column of Aggregation Scheme

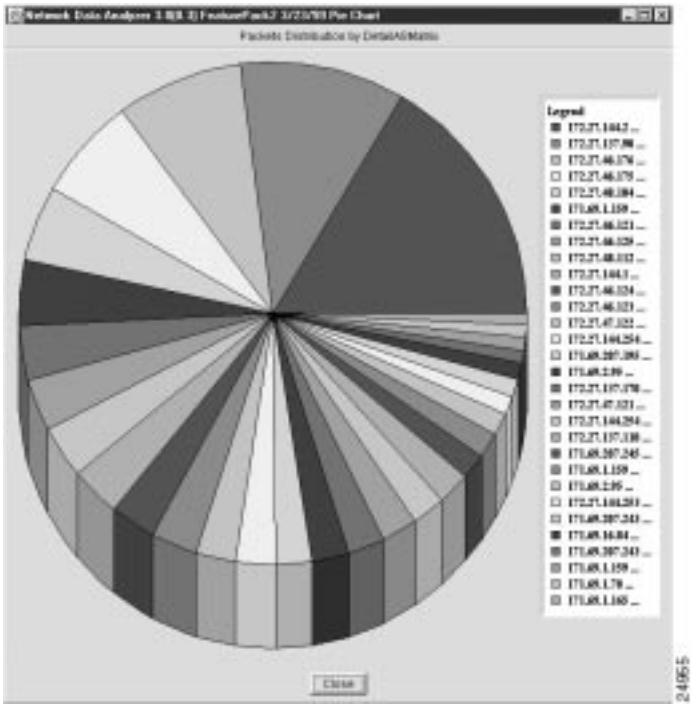

- Method 2—Click the Graph pull-down menu of the Display module window. This pops up a window listing Bar, Pie, or Histogram as options.

  Place the mouse pointer on Pie. A list of choices for the TopN value appears. Move the mouse pointer to the desired value (assumed to be “30” for this step) and click the value to initiate pie chart creation.

  The resulting pie chart (see Figure 3-33) appears in a separate window on your Display module screen.

  To remove the pie chart from the screen, click Close.

Creating a Histogram

For the purpose of creating a histogram, use the data in the DetailInterface aggregation scheme shown in Figure 3-34. Note that the Bytes column in the data array is highlighted, indicating that the data in this column is to be used as the basis for creating the histogram. Note also that you must select a row in the overall NetFlow data array for which you want a histogram representation of traffic data.

Note The Analyzer does not support the display of TMS data as a histogram chart.

Figure 3-34 Typical Aggregation Scheme Data for Creating Histogram Chart

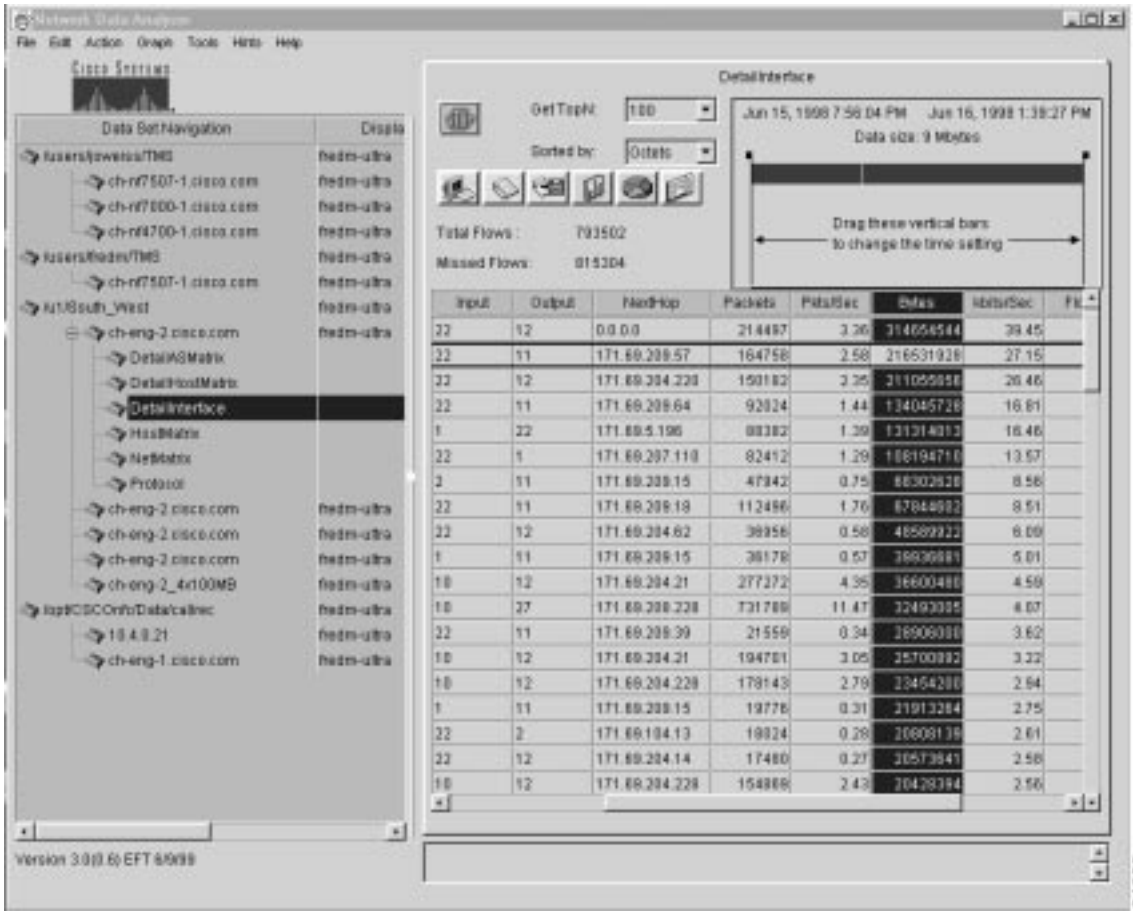

You can use either of two methods in creating a histogram representation of the traffic information in the Bytes column (or any other value column of a NetFlow data aggregation scheme):

- Method 1—Click the Draw Histogram Chart button in the display pane.

  A "Histogram Chart in Progress" pop-up window appears in the Display module window to show the completion percentage of the graphing function.

  The histogram chart shown in Figure 3-35 appears in a separate window.

  To remove the histogram from the screen, click Close.

Figure 3-35 Histogram Chart for Selected Column and Row of Aggregation Scheme

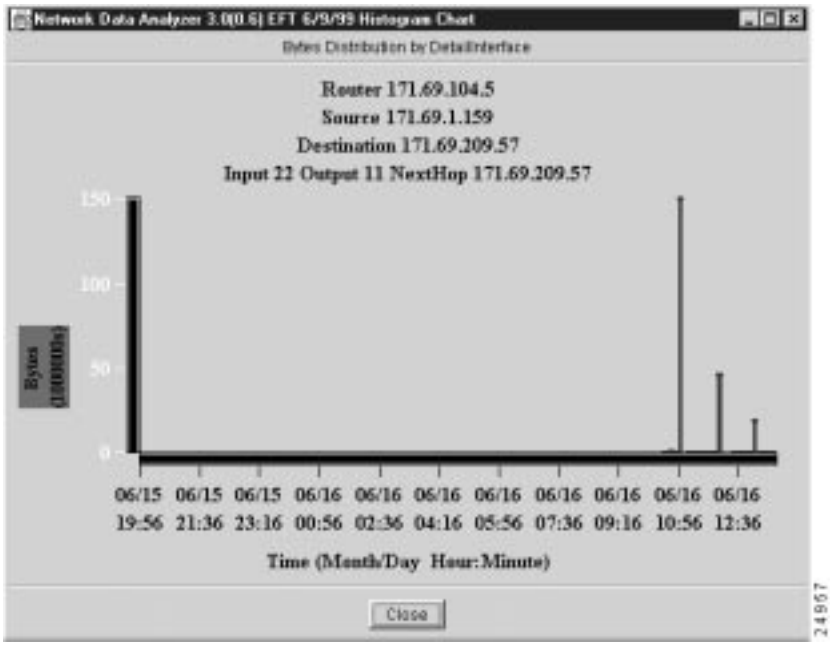

- Method 2—Click the Histogram option of the Graph menu in the Display module window.

  A "Histogram Chart in Progress" pop-up window appears temporarily in the Display module window to show the completion percentage of the graphing function.

  The histogram chart shown in Figure 3-35 appears in a separate window.

  To remove the histogram from the screen, click Close.

Tools Menu Options for Data Exploration

This section describes the following Tools menu facilities for data exploration:

- AS Drill Down Window—See the "Drilling Down on Network Flows" section below.
- Search Window—See the "Searching for Flows by Source and Destination Addresses" section on page 3-57.

Drilling Down on Network Flows

The AS Drill Down Window option enables you to take a closer look at traffic data pertaining to multiple devices involved in traffic flows between source AS and destination AS systems in your network. An AS (autonomous system) is a network, or a collection of networks, operating under a common network administration and routing strategy.

In AS networks, it is often useful to examine traffic data closely for several devices in the network to determine the best way to administer the network and use its resources. The AS Drill Down Window of the Display module provides this capability.

The AS Drill Down Window option of the Tools menu is used in connection with the DetailASMatrix aggregation scheme. Figure 3-36 shows a typical data array for DetailASMatrix aggregation schemes for three selected devices.

Figure 3-36 DetailASMatrix Aggregation Scheme Data for AS Drill Down Operations

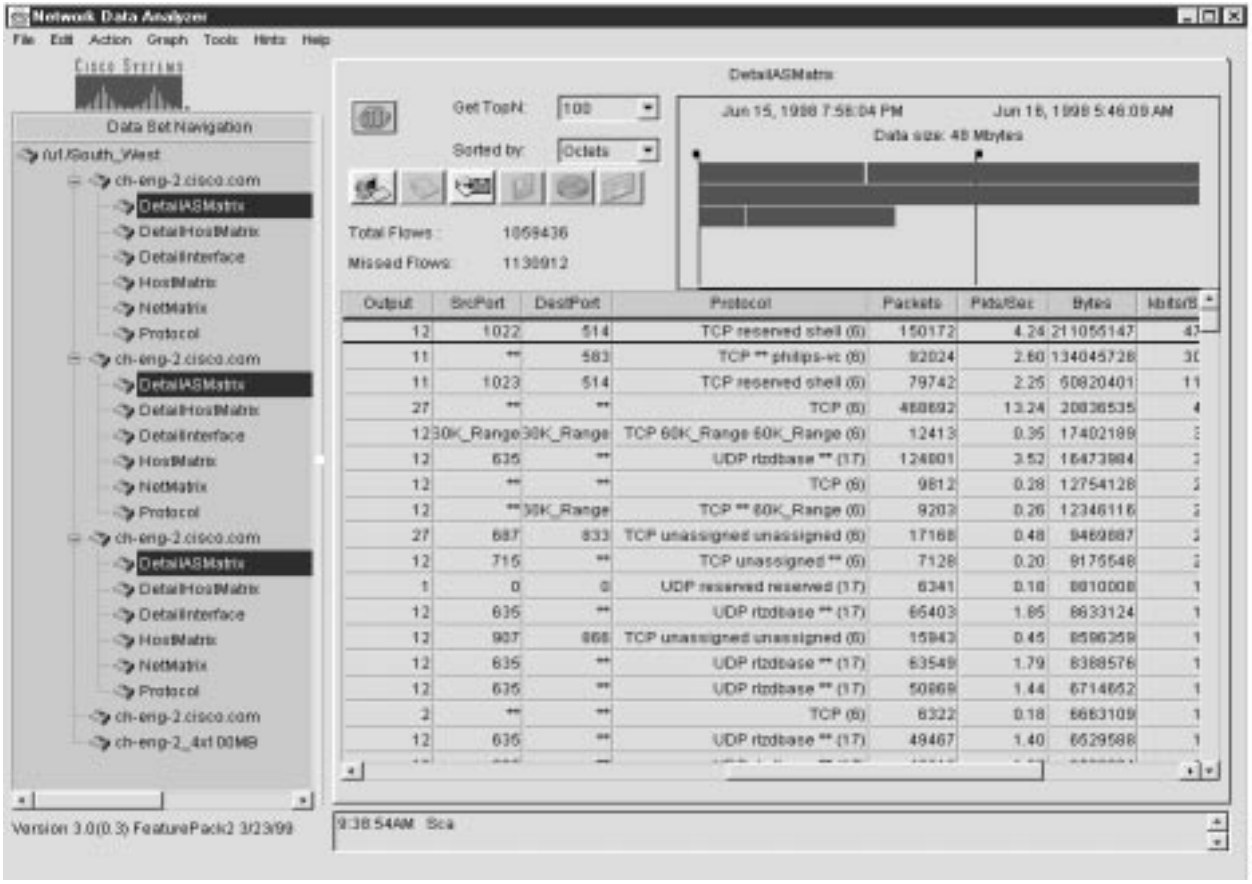

To drill down on NetFlow data pertaining to source and destination AS systems in your network, perform the following steps:

Step 1 Select the AS Drill Down Window option from the Tools menu of the Display module. The AS Drill Down Window appears (Figure 3-37).

The data tree structure shown in the Data Set Navigation pane of Figure 3-36 is propagated into the equivalent area of the AS Drill Down window for device selection purposes in initiating AS drill down functions. In addition, the time line area in the top right portion of the window indicates the ranges of dates and times for which data is available for the selected routers and aggregation schemes.

Step 2 Select those devices that you want to use for AS drill down operations.

Assume, for example, that you want to include the first three devices in the router group named "/u1/South_West," as follows:

- ch-eng-2.cisco.com
- ch-eng-2.cisco.com
- ch-eng-2.cisco.com

While holding down the **Ctrl** key, click each device name to highlight it.

Note Note that the three selected devices have the same name. This circumstance is not unusual, since the same device can be exporting NetFlow data from multiple ports.

Figure 3-37 AS Drill Down Window of the Display Module

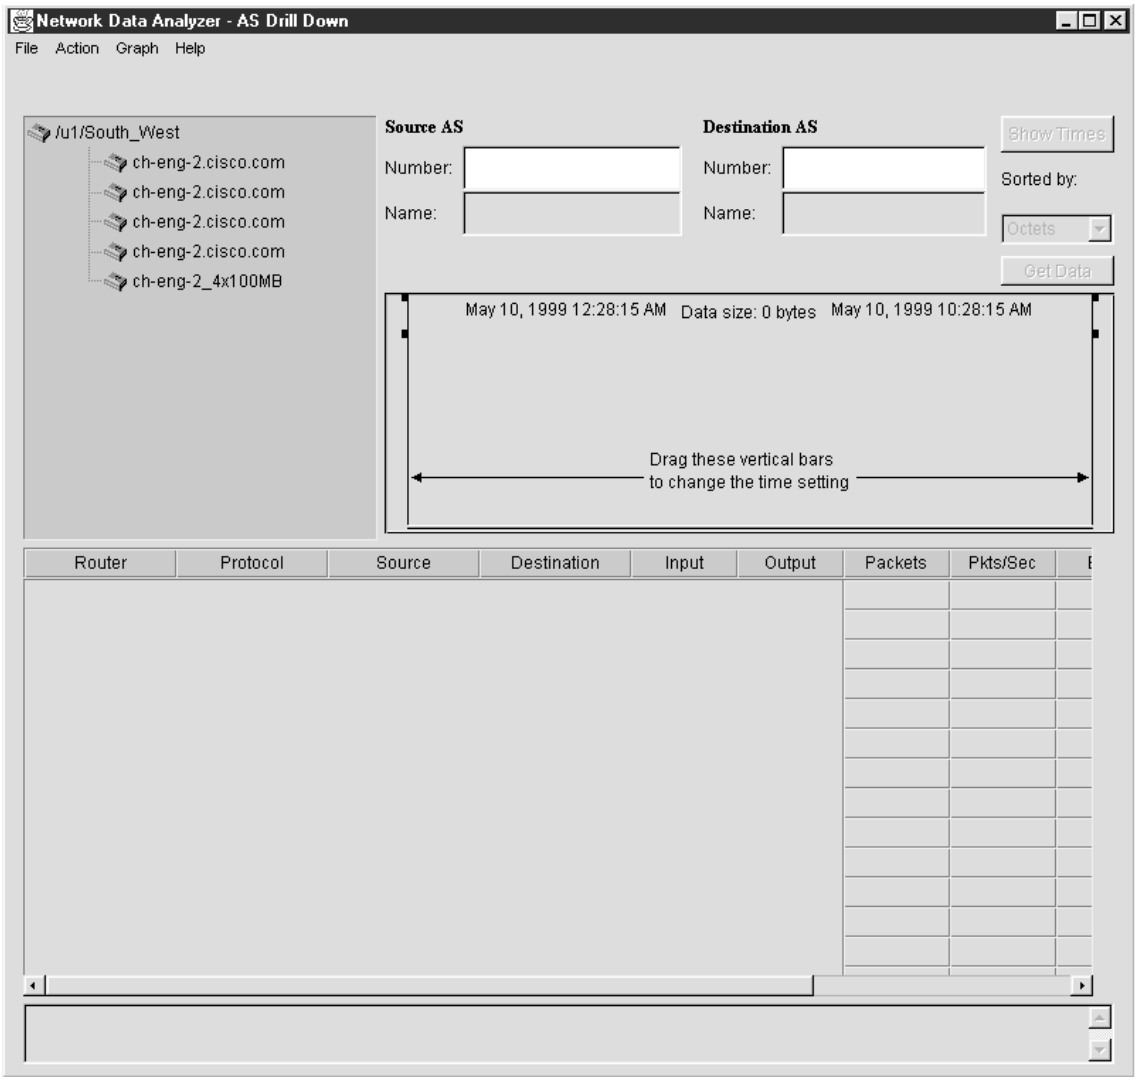

Step 3 Establish the desired parameters for performing the AS drill down operations, as follows:

(a) Enter the Source AS number and the Destination AS number in the appropriate fields of the AS Drill Down window. For this example, assume that "0" is entered in each field as the source and destination AS numbers.

(b) Select the Show Times option under the Action menu of the AS Drill Down window or click the Show Times button.

This action shows the time spans for which aggregation scheme data has been collected for the selected devices.

(c) Position the time slider marks to establish the desired time setting for retrieving traffic data for the three selected devices.

(d) Choose a value for the Sorted by: field. Choices include the following:

- Packets
- Octets
- Flows
- Data key

For this example, choose Octets as the sort key.

(e) Click the Get Data button in the ASDrill Down window to initiate data retrieval.

Alternatively, you can select the Get Data option from the Action menu to initiate data retrieval.

Either action results in the display of appropriate AS data for the selected devices, as shown in Figure 3-38. The display output shows which protocols were used by each selected router interface.

Figure 3-38 Results of Data Retrieval Operations for AS Drill Down Function

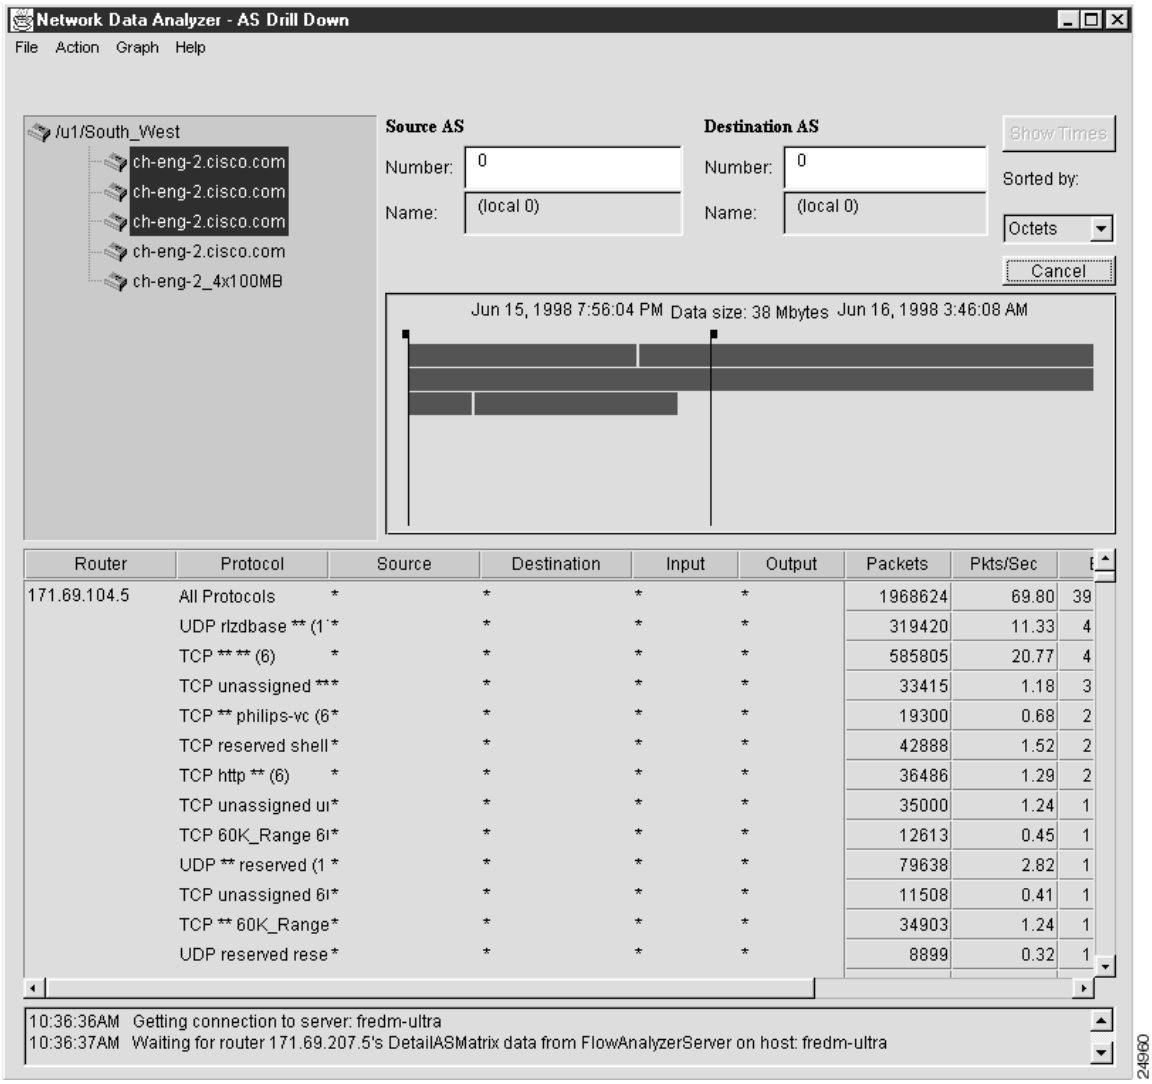

Step 4 To drill down on (get more detail for) the protocols used by a selected device, do the following:

(a) If necessary, use the vertical scroll bar in the right margin of the AS Drill Down window to bring into view the desired data for the device of interest.

(b) If necessary, click and drag the sizing bar in the right margin of the Protocol column header to expand the column enough to bring all information in the column into view.

(c) Select the desired row in the Protocol column for the selected device.

For this step, assume that you want to drill down on the protocol data for the third selected data set path for the device named “ch-eng-2.cisco.com.” The resulting data for this device will be identified in the resulting data array by its IP address 171.69.207.5.

To bring the data for this interface into view, it may be necessary to use the vertical scroll bar in the window to traverse to the appropriate area of the data array.

Assume further that you want to get more detail on the third row in the Protocol column for this device, that is, the row containing the protocol “UDP rtzdbase**(17),” and that you want to use this row as the basis for drill down operations (Figure 3-39).

Figure 3-39 Selected Row of Protocol Column for AS Drill Down Operations

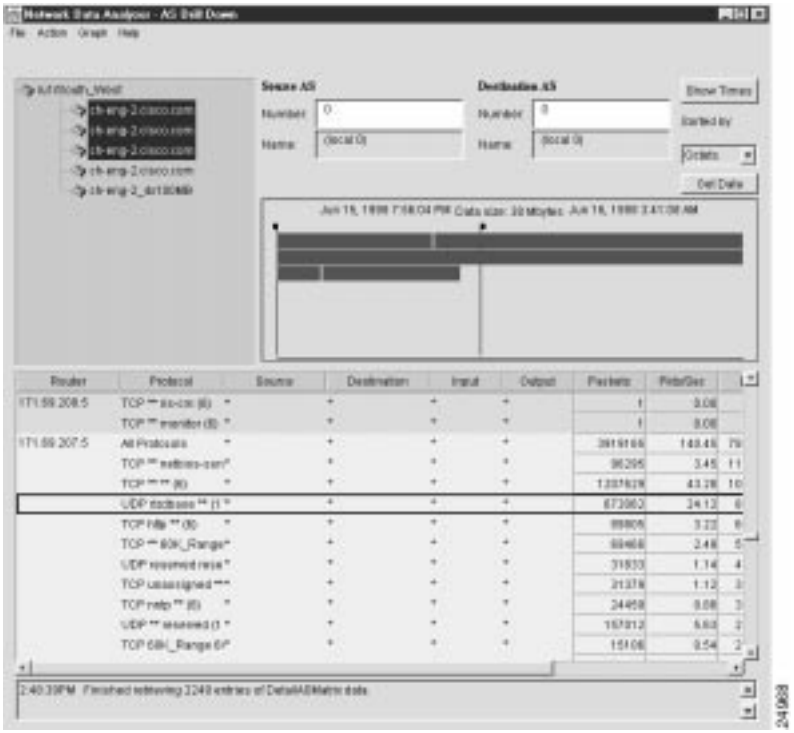

Double-click this row to initiate the display of AS drill down data. As shown in Figure 3-40), detailed traffic flow information is displayed for each individual source destination pair seen on the interface with IP address 171.69.207.5.

Alternatively, you can select the Drill Down on Protocol option of the Action menu in this window to initiate the display of AS data.

Step 5 To collapse (hide) the AS data currently being displayed for the selected row, select the Hide Drill Down Data option from the Action menu in the window.

This action causes the NetFlow data array to revert to its previous form (as shown in Figure 3-39), at which point, you can select any other Protocol row for any other selected device and use that row as the basis for AS drill down operations.

You can perform a series of AS drill down operations for any number of Protocol rows in the data array for a selected device.

Figure 3-40 Result of AS Drill Down Operations

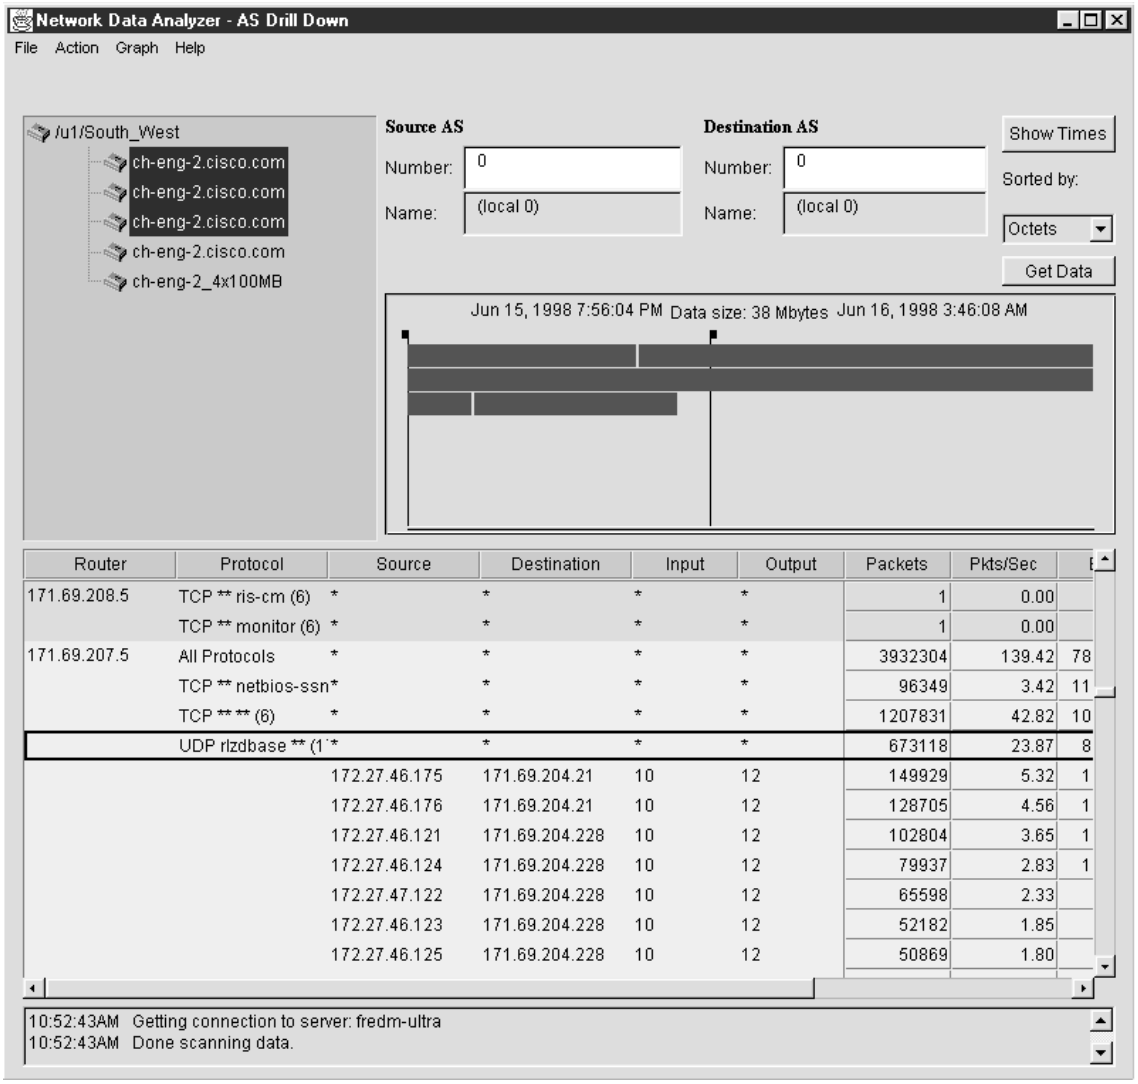

Using Asterisks in the AS Drill Down Window

In the AS Drill Down window, an asterisk appearing in any Protocol row of the traffic data array indicates that additional information relating to the DetailASMatrix aggregation scheme is available for display purposes.

For any such row, you can do either of the following:

- Double-click the row to display additional protocol data for that row.
- Double-click the expanded row to collapse it and return it to its previous state.

By successively double-clicking a series of rows containing an asterisk, you can open each row and view more detail about the specific end-to-end flows with the specific protocol seen on that selected interface.

Effect of Drilling Down on Protocol Row

The effect of drilling down on any row in the Protocol column of an AS Drill Down window that contains asterisks is to expand the following columns of the displayed data aggregation scheme:

- Source and Destination columns—You can expand these columns of an AS Drill Down window in unison to display the IP addresses of all of the communicating source and destination devices in the network whose traffic flows through the selected exporting device.
- Input and Output columns—The integers appearing in Input and Output columns correspond to the MIB ifIndex values of the physical interfaces of the communicating network devices.

Searching for Flows by Source and Destination Addresses

The Search option of the Tool menu enables you to search for traffic flows in the following ways:

- "Searching for IP-to-IP Transactions" section on page 3-60—Flows between one IP address and another IP address
- "Searching for IP-to-Subnet Transactions" section on page 3-61—Flows between one IP address and specified subnet addresses
- "Searching for Subnet-to-Subnet Transactions" section on page 3-62—Flows between a specified subnet address and another specified subnet address
- "Searching for IP "Away From" Subnet Transactions" section on page 3-63—Flows that one IP address sends anywhere in the network, except for another IP address or specified subnet address

Procedures for initiating searches for the above types of traffic flows appear in the following sections.

Preparing for Search Operations

Use the Search Window option of the Tools menu only in connection with the DetailASMatrix aggregation scheme. Before you can perform search operations involving this aggregation scheme, the following conditions must exist:

- Ensure that traffic data for the DetailASMatrix aggregation scheme has been collected from one or more NetFlow export-enabled devices in your network.
- Ensure that the tree structure in the Data Set Navigation pane refers to one or more NetFlow data exporting devices in your network.

Figure 3-41 shows a sample data tree structure that meets these requirements.

When you select the Search Window option of the Tools menu, a separate Search window appears on the Display module screen (see Figure 3-42). Note that the current contents of the Data Set Navigation pane are propagated into the pane labelled "Using these Routers" for device selection purposes.

Figure 3-41 Typical Data Tree Structure for Search Operations

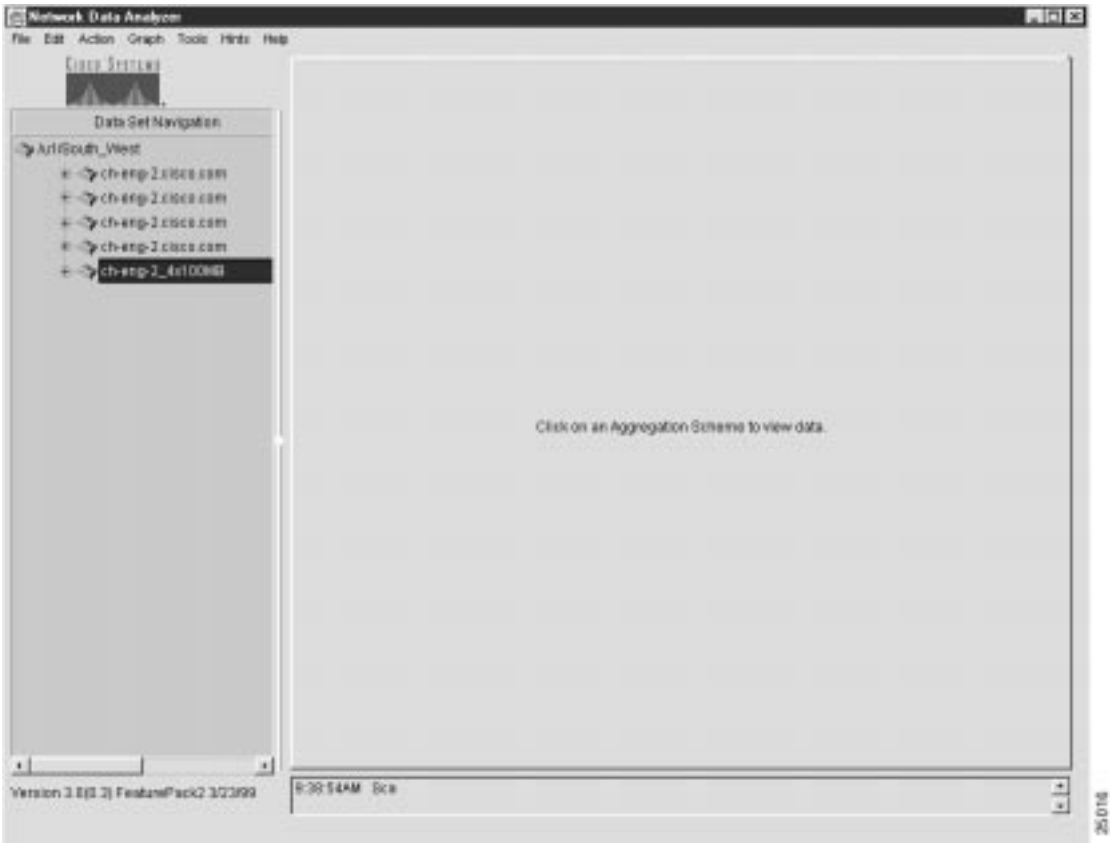

Figure 3-42 Search Window for DetailASMatrix Aggregation Scheme

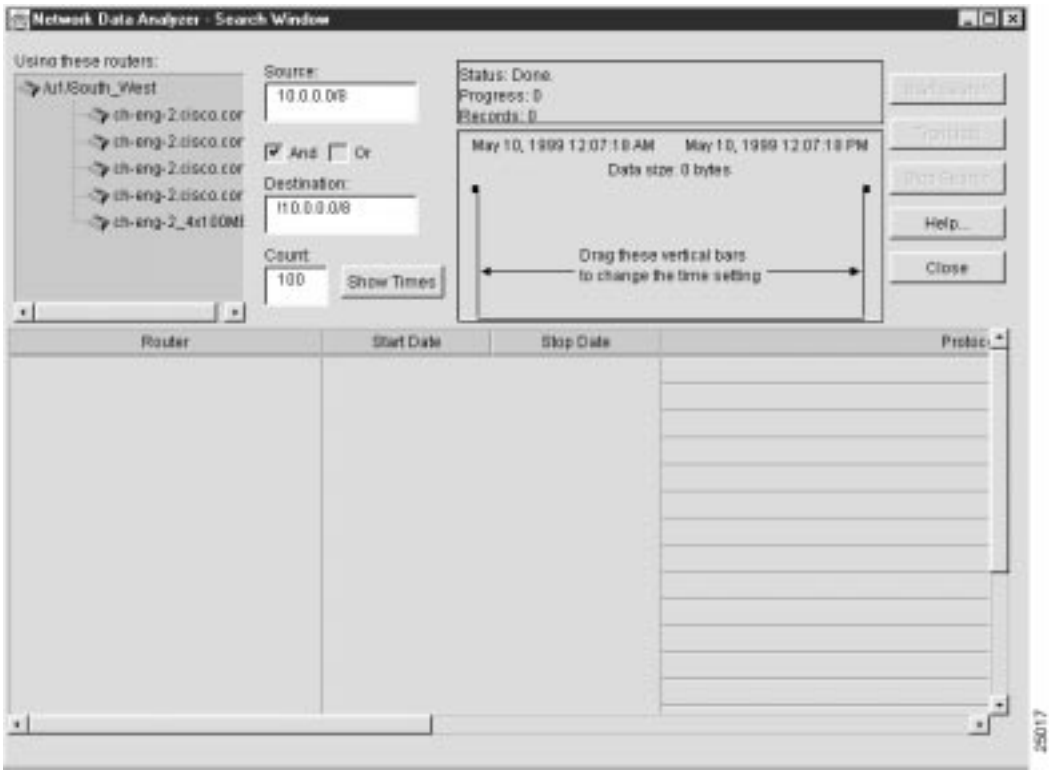

Generalized Outline of Search Procedure

When you access the Search Window by invoking the Search Window option of the Tools menu, you can initiate search operations by performing the following steps:

Step 1 Select the device (or devices) of interest in the data tree structure of the Search Window.

In choosing devices that you want to apply for search purposes, you have the following options:

- Option 1—Choose all the devices in the data tree structure

  When you click the name of the router group (/u1/South_West, for example) in the "Using these routers" pane of the Search Window (see Figure 3-42), you can select all the devices listed as members of the group collectively.

- Option 2—Select devices individually in the data tree structure by holding down the **Ctrl** key as you click each desired device name.

Step 2 Enter the appropriate IP address information for search operations in the Source and Destination fields of the Search Window.

When you enter IP addresses or subnet addresses in either the Source field or the Destination field of the Search Window, use the format shown in the following example in specify the range of IP addresses:

```
192.69.0.0/16
```

Step 3 Click the AND box or the OR box, as appropriate, to define the desired boolean operator for the intended search operation.

Step 4 Click the Show Times button.

Step 5 Position the time slider marks, as appropriate, to define the time horizon for which DetailASMatrix aggregation scheme data is to be retrieved and processed during search operations.

Step 6 Click the Start Search button.

The different searches that you can perform using this generalized procedure are described in the following sections.

Searching for IP-to-IP Transactions

This search operation looks for traffic flows that have occurred between a specified source device and a specified destination device.

To search for traffic flows between one IP address and another IP address, perform the following steps:

Step 1 In the Source field, enter the IP address of the source device that originated the flows (171.69.210.178, for example).

Step 2 In the Destination field, enter the IP address of the destination device that received the flows (161.44.128.38, for example).

Step 3 Check the AND box.

Step 4 Click the Show Times button, adjust the time slider marks, as desired, and click the Start Search button.

This procedure finds the traffic flows that originated from the source device (171.69.210.178) and that were received by the destination device (161.44.128.38).

Sample output from an IP-to-IP search operation is shown in Figure 3-43.

Figure 3-43 Sample Output from IP-to-IP Search Operation

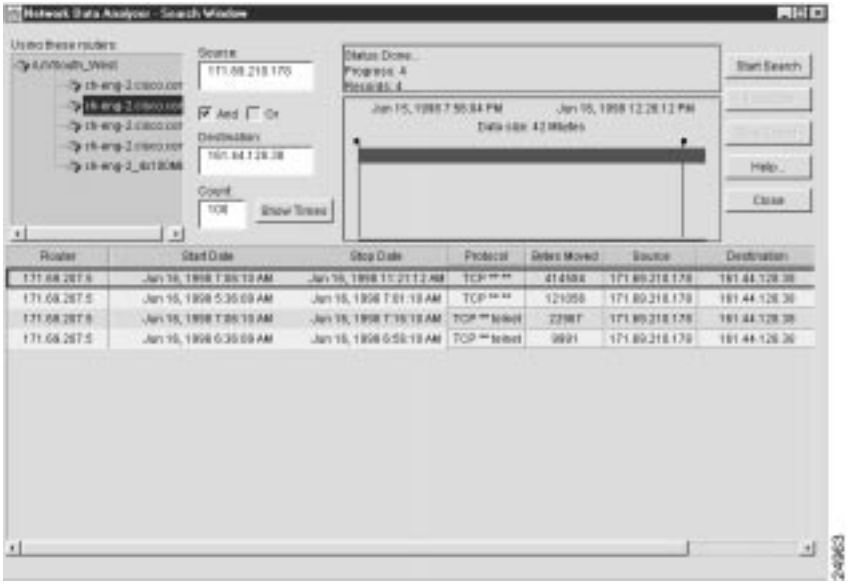

Note The Search window contains a Translate button in the upper right corner of the window. This button serves the same function as the equivalent button or menu item in other Display module windows, namely, to translate host IP addresses into DNS names.

Searching for IP-to-Subnet Transactions

This search operation looks for traffic flows that have occurred between a specified source device and destination devices that have IP addresses within a specified range on a subnet.

To search for flows between a specified source device and a range of subnet addresses, perform the following steps:

Step 1 In the Source field, enter the IP address of the source device that originated the flows (171.69.210.178, for example).

Step 2 In the Destination field, enter the subnet address of the destination devices that received the flows (161.44.0.0/16, for example).

Step 3 Check the AND box.

Step 4 Click the Show Times button, adjust the time slider marks, as desired, and click the Start Search button.

This procedure finds the flows that originated from the source device with the IP address 171.69.210.178 and that terminated at destination devices having IP addresses in the range 161.44.0.0 to 164.44.255.255.

Figure 3-44 shows sample output from the IP-to-subnet search operation above.

Figure 3-44 Sample Output from IP-to-Subnet Search Operation

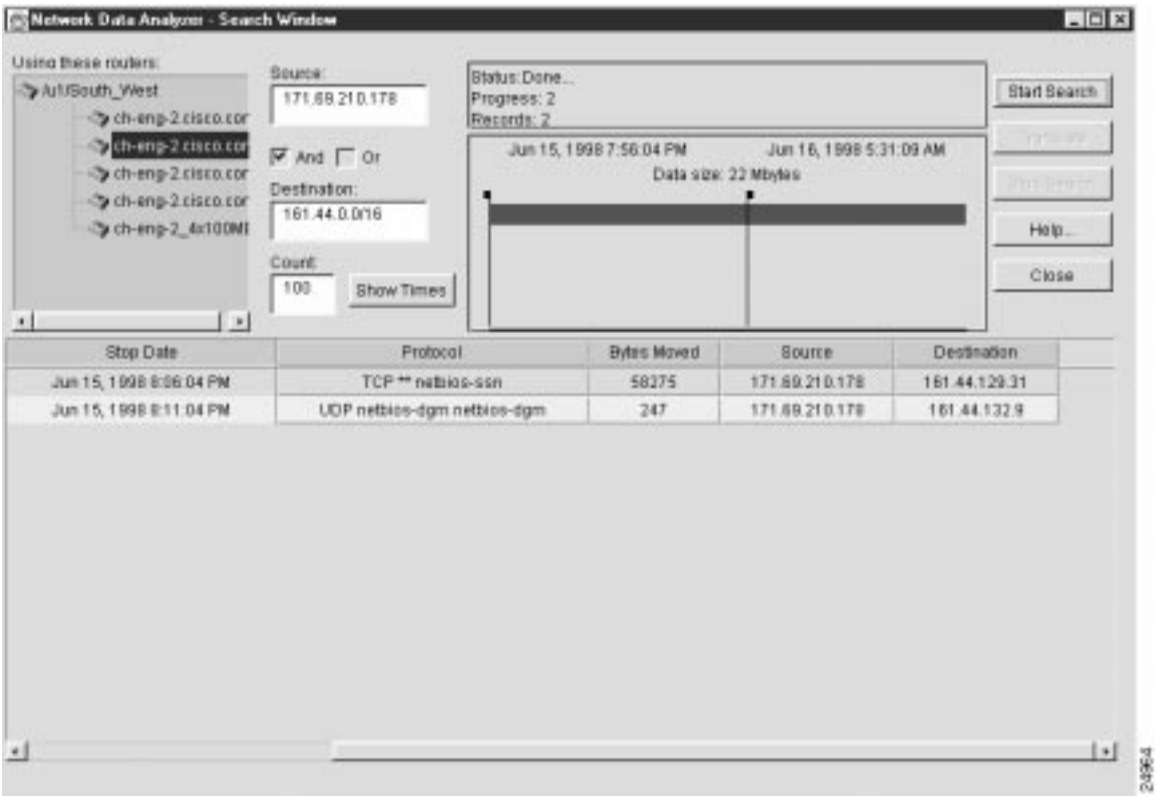

Note that you can swap the contents of the Source and Destination fields to accomplish the reverse of the search operation described above. This procedure is described in the following section.

To search for flows between a range of subnet addresses and a specified IP address, perform the following steps:

Step 1 In the Source field, enter the subnet address of the source devices that originated the flows (161.44.0.0/16, for example).

Step 2 In the Destination field, enter the IP address of the destination device that received the flows (171.69.210.178, for example).

Step 3 Check the AND box.

Step 4 Click the Show Times button, adjust the time slider marks, as desired, and click the Start Search button.

This procedure finds the flows that originated from the source devices having IP addresses in the range from 161.44.0.0 to 161.44.255.255 and that were received by the destination device at IP address 171.69.210.178.

Searching for Subnet-to-Subnet Transactions

This search operation looks for traffic flows that have occurred between source devices having IP addresses within a specified range and destination devices having IP addresses within a specified range.

To search for flows between one subnet address and another subnet address, perform the following steps:

Step 1 In the Source field, enter the subnet address of the source devices that originated the flows (171.69.0.0/16, for example).

Step 2 In the Destination field, enter the subnet address of the destination devices that received the flows (161.44.0.0/16, for example).

Step 3 Check the AND box.

Step 4 Click the Show Times button, adjust the time slider marks, as desired, and click the Start Search button.

This procedure finds the flows that originated from the source devices having an IP address in the range from 171.69.0.0 to 171.69.255.255 and that were received by the destination devices having an IP address in the range from 161.44.0.0 to 161.44.255.255.

Figure 3-45 shows sample output from the subnet-to-subnet search procedure outlined above.

Figure 3-45 Sample Output from Subnet-to-Subnet Search Operation

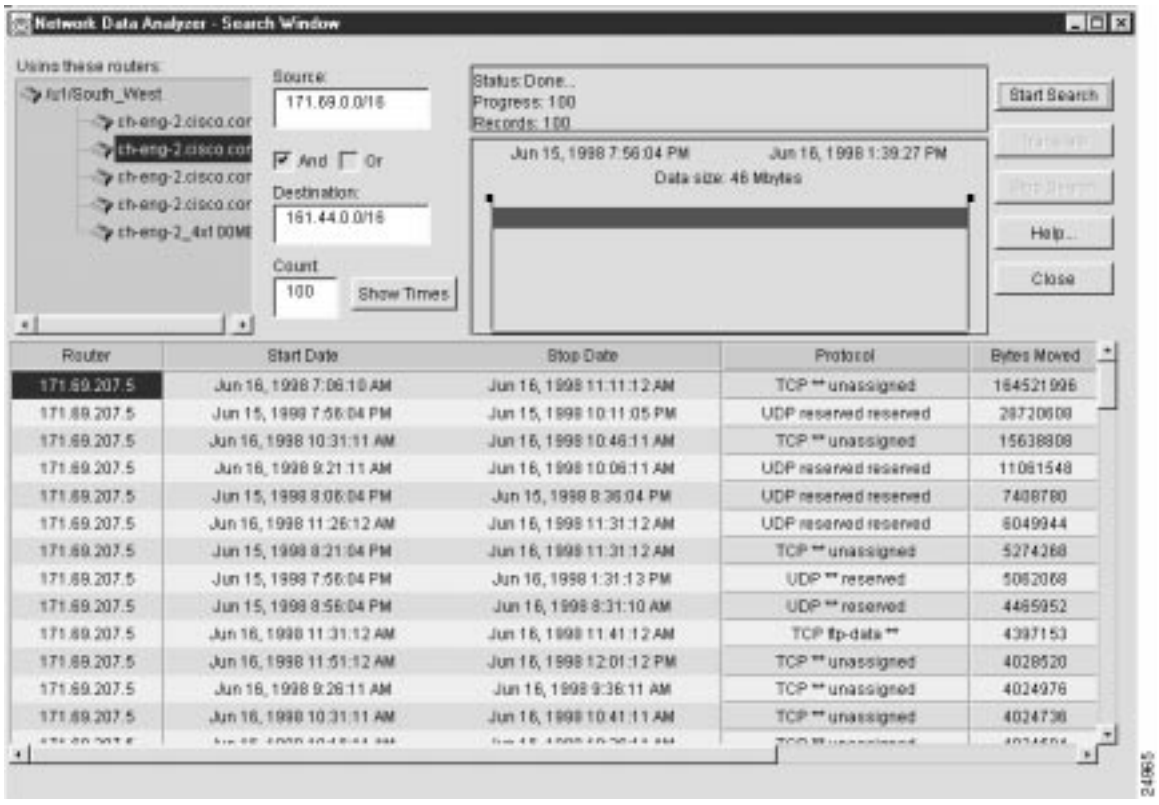

Searching for IP “Away From” Subnet Transactions

This search operation looks for traffic flows that have occurred between a specified source device and destination devices other than those having IP addresses within a specified subnet’s range.

To search for flows that a specified device directs to other devices in the network, except to certain specified devices, perform the following steps:

Step 1 In the Source field, enter the IP address of the source device that originated the flows (171.69.210.178, for example).

Step 2 In the Destination field, enter the ! character, followed by a subnet address range (161.44.0.0/16, for example).

The exclamation point (!) is a "not" operator that means "any but those" addresses specified in the subnet address. The effect of this operator is to limit the scope of the search operation.

Step 3 Check the AND box.

Step 4 Click the Show Times button, adjust the time slider marks, as desired, and click the Start Search button.

This procedure finds the flows that originated from the specified device (171.69.210.178) and terminated at devices other than those having an IP address in the specified range (161.44.0.0 to 161.44.255.255).

Figure 3-46 shows sample output from the search procedure outlined above.

Figure 3-46 Sample Output of IP "Away From" Subnet Search Operation

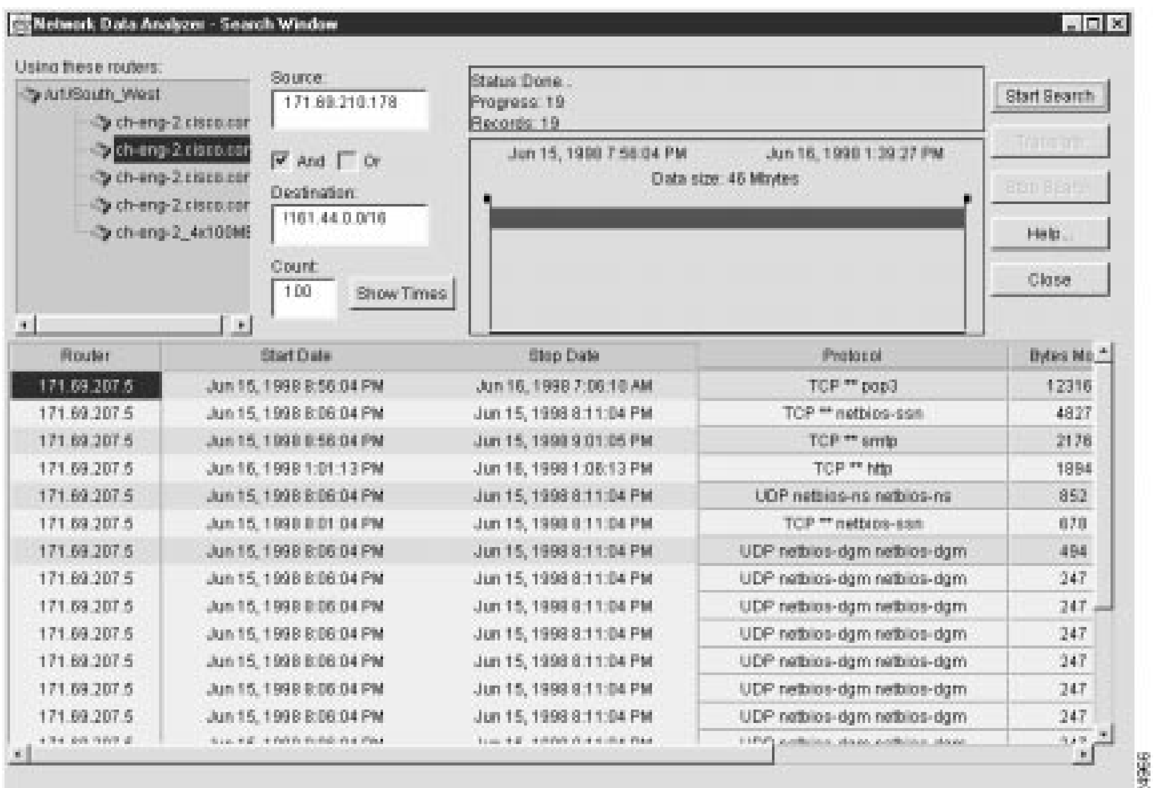

Note that you can swap the contents of the Source and Destination fields to accomplish the reverse of the search operation described above. This reverse procedure is described in the following section.

To search for flows that originate anywhere in the network (except from devices having an IP address in a specified range) and that terminate with a specified device, perform the following steps:

Step 1 In the Source field, enter the ! operator, followed by a subnet address (161.44.0.0/16, for example).

The exclamation point (!) is a “not” operator that means “any but those” addresses specified in the subnet address. The effect of this operator is to limit the scope of the search operation.

Step 2 In the Destination field, enter the IP address of the destination device that received the flows (171.69.210.178, for example).

Step 3 Check the AND box.

Step 4 Click the Show Times button, adjust the time slider marks, as desired, and click the Start Search button.

This procedure finds all the traffic flows that originated from anywhere in the network (except for those devices having an IP address in the range from 161.44.0.0 to 161.44.255.255) and that were received by the destination device at IP address 171.69.210.178.

Tools Menu Options for Data Collection

This section describes the following Tools menu facilities for data collection:

- TMS Collection Control—See the “Controlling TMS Data Collections” section below.
- NetFlow Collection Control—See the “Controlling NetFlow Data Collections” section on page 3-72.
- Router Configuration—See the “Configuring Routers for Data Export” section on page 3-103.

Controlling TMS Data Collections

This section describes the facilities provided by the Analyzer for controlling traffic matrix statistics (TMS) data collections.

Introduction to the Traffic Matrix Statistics Control Window

Depending on whether or not TMS collections have already been defined, either of two versions of the Traffic Matrix Statistics Control window appears when you select the TMS Collection Control option of the Tools menu:

- TMS collections found—In this case, one or more TMS collections have already been defined. Thus, when you select the TMS Collection Control option of the Tools menu, the Traffic Matrix Statistics Control window shown in Figure 3-47 appears.

Figure 3-47 Traffic Matrix Statistics Control Window

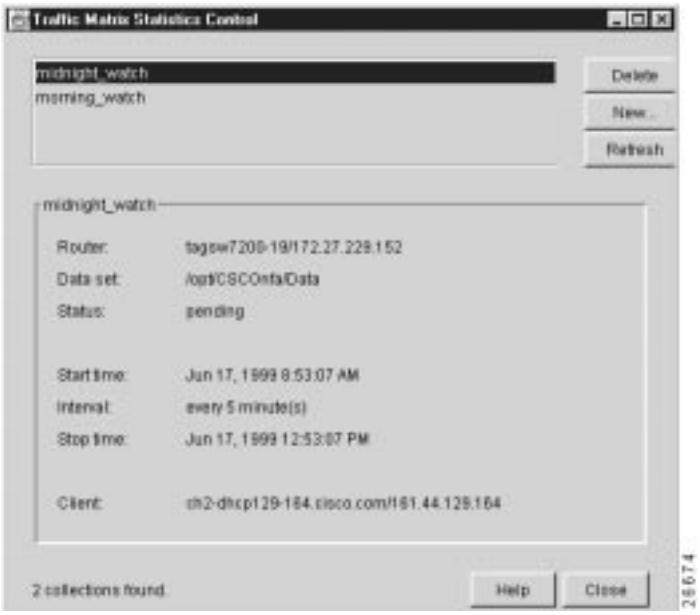

The bottom left corner of this window indicates that two TMS collections have been found. These two collections are identified by name in the rectangular area at the top of the window as "midnight_watch" and "morning_watch."

The first collection name listed in the top of the window ("midnight_watch") is highlighted by default when the Traffic Matrix Statistics Control window is displayed. In addition, the bottom portion of the window displays the associated TMS collection control parameters for the collection named "midnight-watch."

There are as many TMS collection names listed in the top of the Traffic Matrix Statistics Control window as there are existing defined TMS collections. If the list of named TMS collections in the window exceeds the space available, a vertical scroll bar appears in the window, enabling you to traverse to any desired collection name and select if for collection control purposes.

The "Defining Parameters for a New TMS Collection" section on page 3-68 describes how to define a new TMS collection.

- No TMS collections found—In this case, no TMS collections have been previously defined, causing a "blank" Traffic Matrix Statistics Control window to appear (see Figure 3-48) when you select the TMS Collection Control option of the Tools menu.

Figure 3-48 "Blank" Traffic Matrix Statistics Control Window

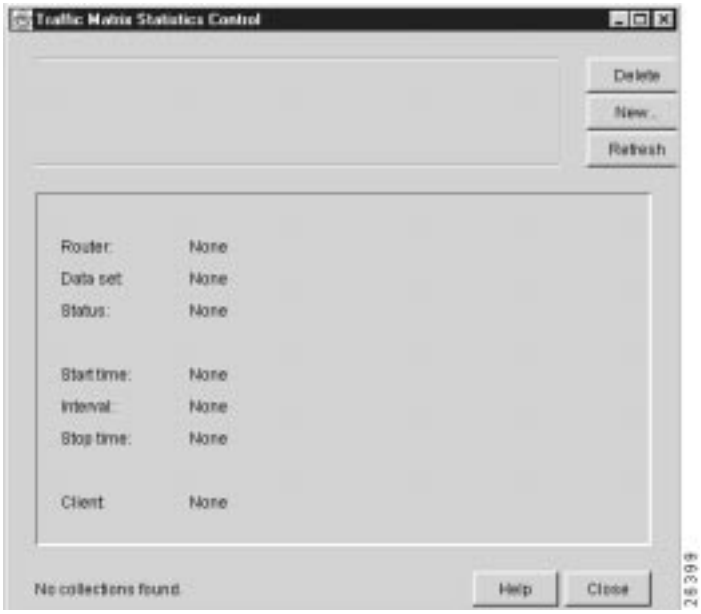

Note that the lower left corner of the Traffic Matrix Statistics control window contains the notation "no collections found." Note further that no TMS collections are listed in the rectangular collection name area at the top of the window.

Given these conditions, you must click the New button in the Traffic Matrix Statistics Control window to define a new TMS collection, as described in the following section.

Specifying a Configuration Files Directory for a New TMS Collection

For purposes of this section, it is assumed that no TMS collections have been defined. Thus, when you click the New button in the Traffic Matrix Statistics Control window, the Configuration Files Directory window shown in Figure 3-49 appears, prompting you to enter the name of a directory on a host UtilityServer that contains configuration files for the routers and switches in your network.

Figure 3-49 Configuration Files Directory Window

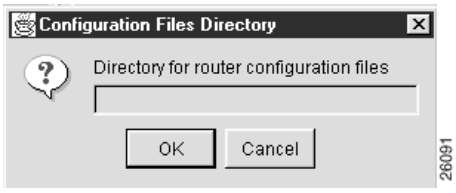

Typically, the directory path name that you enter into the Configuration Files Directory window points to the directory "/tftpboot/configs/" created previously on a host UtilityServer for storing router configuration files.

There must be a configuration file in this directory for each TMS export-capable device that you intend to control by means of the Traffic Matrix Statistics Control window. You create router configuration files at installation time when you configure devices to operate in your network. Thus, for each new device so configured, you log on to the device and copy its running configuration file into the /tftpboot/configs/ directory.

The configuration file directory provides the names, interfaces, and passwords of network devices, enabling you to log on to any TMS export-capable device in the network and to configure that device for TMS collections by means of the Traffic Matrix Statistics Control window.

Typically, you store the router configuration files on a designated UtilityServer host in the network. However, you can store such files on any NFS-mounted volume in the network that is accessible to the UtilityServer. In other words, the directory path name that you enter in the Configuration Files Directory window can point to a directory on any UNIX platform in the network that is reachable by the UtilityServer.

After you enter a directory path name in the Configuration Files Directory window, click OK. If the directory path name that you enter is valid, the Traffic Matrix Statistics Control window shown in Figure 3-47 appears. Select any named collection listed in the rectangular TMS collection name area at the top of this window as the target for configuring the new TMS collection parameters. See the next section for a description of how you define the parameters for a new TMS collection.

If you decide to abandon the directory specification task, click Cancel in the Configuration Files Directory window to return to the Traffic Matrix Statistics Control window.

If the directory path name that you enter in the Configuration Files Directory window contains no configuration files, or if you specify the directory path name incorrectly, the following Router Config Files window appears to so indicate.

Defining Parameters for a New TMS Collection

You define a new TMS collection by setting specific parameters in the Traffic Matrix Statistics Control window.

When you click the New button in the Traffic Matrix Statistics Control window and a valid router configuration files directory exists on a designated UtilityServer host in the network, the Traffic Matrix Statistics Control window shown in Figure 3-50 appears. This window incorporates a New Collection panel that enables you to define a new TMS collection process.

Figure 3-50 New Collection Panel for Defining Collection Control Parameters

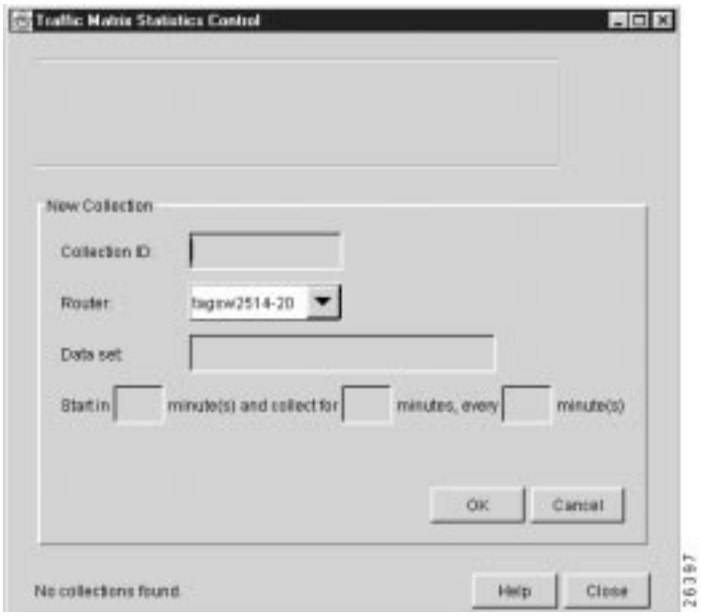

The New Collection panel provides facilities for defining the following TMS collection parameters:

- Collection ID field—Use this field to compose an alphanumeric name of your choosing (of any length without embedded spaces) by which the TMS collection process on the selected router (see next bullet) is to be identified.

  For example, “midnight_watch” (or any other meaningful name of your choice) can be entered as the name of the TMS collection process.

- Router—Use this single-selection, pull-down box to select the name of a network device for which you want to define a new TMS collection. This parameter identifies the device that is to be polled for TMS data.

  The devices listed in this pull-down box mirror those listed in the specified configuration files directory (see the “Specifying a Configuration Files Directory for a New TMS Collection” section on page 3-67).

  You can select any one of the devices listed in the Router pull-down selection box as the target device for the new TMS collection definition task. Assume that the device named “tagsw7200-19” has been selected for this purpose.

- Data Set—Use this field to specify the data set path where the exported TMS data is to be stored.

  The data set path that you specify in this field can point to any NFS-mounted volume on a UNIX platform in the network that is accessible to the UtilityServer.

  For example, a typical data set path that you can specify in this field is “/opt/CSCOnfa/Data.”

  If you specify a relative path name in this field, such as “tms,” rather than an absolute path name, the path name “/opt/CSCOnfa/Data/tms” is assumed.

- “Start in”—Use this field to specify how much time (in minutes) is to elapse before the TMS collection process begins. This parameter specifies the time to begin polling the selected router for TMS data.

For example, you could specify a start time that is to take effect in 60 minutes (or any other increment of minutes that is appropriate to your TMS collection requirements).

- "Collect for"—Use this field to specify the overall duration (in minutes) of the TMS collection process. This parameter specifies the frequency at which the UtilityServer requests TMS data from the selected router.

  For example, you could specify a duration of 240 minutes (or any other increment of minutes that is appropriate to your TMS collection requirements).

- "Every"—Use this field to specify (in minutes) how often "snapshots" of the traffic counters in the selected router are to be exported to the designated TMS data repository (see the Data Set field above).

  For example, you could specify that the current contents of the traffic counters in the TMS exporting device are to be exported every 5 minutes (or any other increment of minutes that is appropriate to your TMS data collection requirements).

  If you specify values in the "collect for" and "every" field that are illogical, such as collect for 5 minutes every 10 minutes, the analyzer will warn you about this (or any other such) inconsistency.

Click Cancel if you wish to abandon the TMS collection definition task. Doing so clears the fields of the New Collection panel and returns you to the Traffic Matrix Statistics Control window.

The sample data entered as described above is reflected in the completed New Collection panel shown in Figure 3-51.

Figure 3-51 Completed New Collection Panel

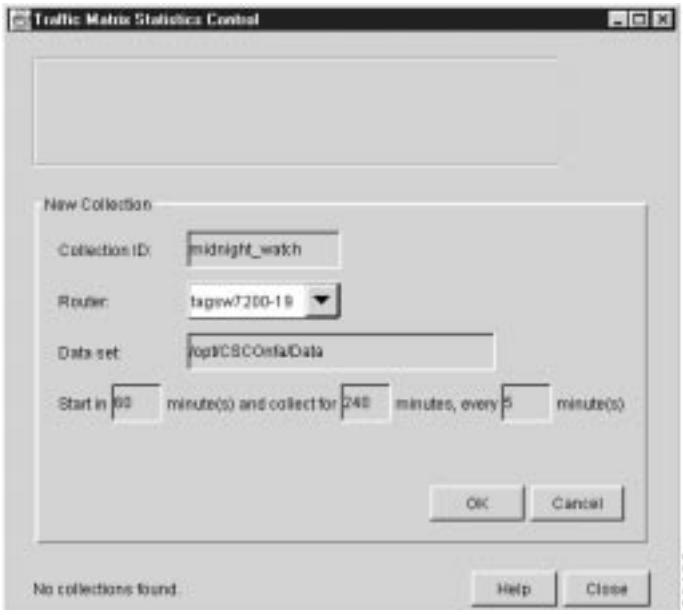

When you click OK in the New Collection panel after specifying the TMS collection control parameters as described in the preceding section, the Traffic Matrix Statistics Control window takes the form shown in Figure 3-52. Note that the new collection name "midnight_watch" now appears at the top left corner of the window.

Figure 3-52 Result of New TMS Collection Definition Task

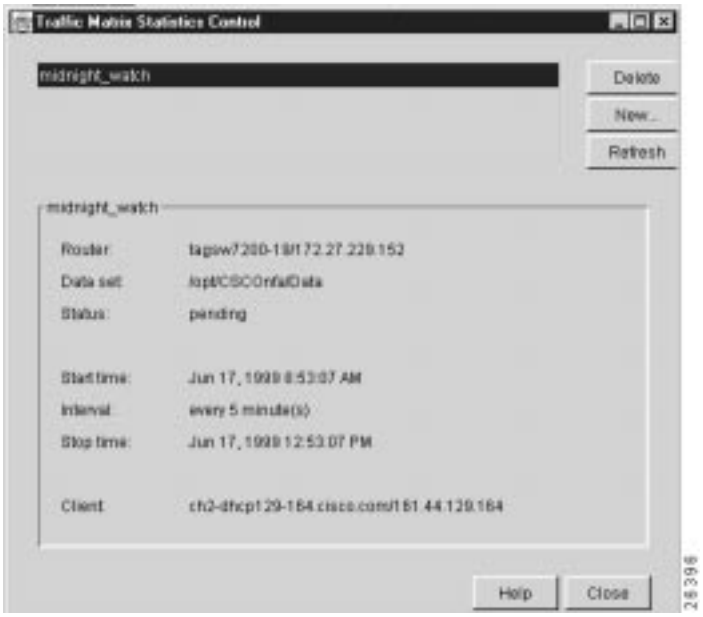

Manipulating a TMS Collection Definition

You can define any number of TMS collections in the manner described in the preceding section, but you can specify only one TMS collection for a given device.

Once TMS collections are defined and listed in the Traffic Matrix Statistics Control window, you can:

- Delete a selected TMS collection—In the collection ID rectangle at the top of the Traffic Matrix Statistics Control window, click the name of the collection that you want to delete to highlight it; then click Delete.

  This collection ID area of the Traffic Matrix Statistics Control window becomes a scrollable pane if more collection names exist than can be listed within the available space in the window.

- Define a new TMS collection—Click New to define a new TMS collection.
- Update all TMS collections—Click Refresh to update the Traffic Matrix Statistics Control window with the latest collection control information for all of the named collections listed in the collection ID rectangle.
- Select a listed TMS collection—When you select a different collection name in the collection ID rectangle, all of the parameters associated with that named collection are updated and displayed in the Traffic Matrix Statistics Control window.

Controlling NetFlow Data Collections

This section describes the configuration and control facilities provided by the NetFlow Collection Control window, which is accessed when you select the NetFlow Collection Control option of the Tools menu.

To collect NetFlow data successfully:

- Routers must be configured to send NetFlow data to a specified FlowCollector.
- A FlowCollector must be configured to receive data from one or more routers.

Therefore, if no FlowCollectors have been defined previously in the UtilityServer's NFCCC.txt file when you select the NetFlow Collection Control option, you must satisfy the requirements outlined in the following sections before you can gain access to the collection control facilities of the NetFlow Collection Control window.

FlowCollector Startup Screens of the Display Module

This section describes a series of screens that the Display module presents to you in the event that no FlowCollectors have been defined previously in the UtilityServer's NFCCC.txt file. If such is the case, the following Network Data Analyzer Start Up dialog box appears.

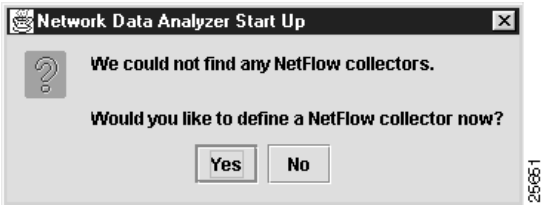

If you click Yes in response to the above message, the following New Collector Dialog box appears, enabling you to enter the name of a FlowCollector to which you want to connect.

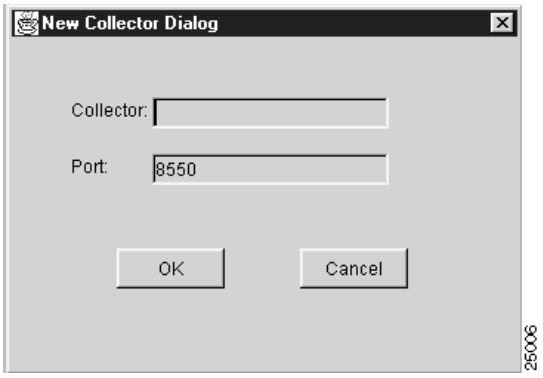

After you enter the FlowCollector name and click OK, the NetFlow Collection Control window appears, displaying the name of the desired FlowCollector in the Collector list of the window.

If you click Cancel in the New Collector Dialog window, the NetFlow Collection Control window appears without displaying the name of a FlowCollector in the Collector list of the window. In this case, if you close the NetFlow Collection Control window and again select the NetFlow Collection

Control option of the Tools menu in another attempt to connect to a FlowCollector, the following error message appears to remind you that no FlowCollectors have been defined in the UtilityServer's NFCCC.txt file.

If you click No in the Network Data Analyzer Start Up window shown at the beginning of this section, the window again appears with a new query: "Would you like to configure routers to export NetFlow Data?"

The "Configuring Routers for Data Export" section on page 3-103 describes how you configure routers for the export of NetFlow or TMS data.

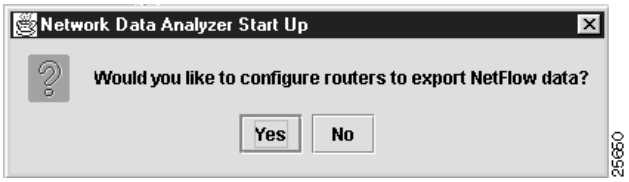

If you click Yes in response to the above configuration query, the following Configuration Files Directory pop-up appears, enabling you to enter the path name of the directory containing the router configuration files on the UtilityServer host.

After entering the path name, click OK.

If you click Cancel in the Configuration Files Directory window above, you will be returned to the main Display module window.

If you click No in response to the "Would you like to configure routers to export Netflow data?" query, the Network Data Analyzer Start Up window shown below appears, querying you: "Would you like to view NetFlow Data?"

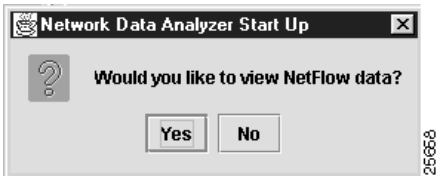

If you answer Yes to the query “Would you like to view NetFlow data?,” the following New Data Location pop-up window appears, enabling you to enter the directory path name (/u1/South_West, for example) that contains the desired NetFlow data.

After you enter the directory path name in the Data Location field, click OK to proceed.

If you click Defaults, the default DisplayServer and port number are set. Also, the current contents of the Data Location field are cleared so that you can enter a desired directory path name.

If you Click Cancel, the directory path name operation is abandoned.

If you answer No to the query “Would you like to view NetFlow data?,” you are returned to the main Display module window without satisfying any of the prerequisites for viewing NetFlow data.

User Name and Password Authentication Process

This section describes the FlowCollector files and Analyzer facilities that come into play during the user name and password authentication process for a FlowCollector. A successful authentication process enables you to gain access to desired NetFlow data stored on a FlowCollector.

FlowCollector Startup Files

The following two files in the FlowCollector’s /opt/CSCOnfc/config directory are of consequence during FlowCollector startup and pertain to the user name and password authentication process:

- nfcd.config file—Contains a list of the programs that can be launched by the FlowCollector Daemon (NFCD). NFCD runs continuously on a FlowCollector host to start and stop programs listed in the nfcd.config file.

  By means of the PROGRAMFLAGS variable of the FlowCollector Gateway (NFCGW) application listed in the FlowCollector’s /opt/CSCOnfc/config/nfcd.config file, you can configure a FlowCollector to: a) perform user authentication on startup, or b) bypass user authentication completely.

  The PROGRAMFLAGS line for NFCGW takes a single parameter, “-i” (dash, lower case i), the absence or presence of which has the effect described below during FlowCollector startup:

  — If -i *is not specified* (the default condition), the FlowCollector enforces user name and password authentication.

    In this case, you are prompted to supply a valid user name and password to the FlowCollector by means of the User Name and Password Dialog window described in the “User Name and Password Authentication Dialog Box” section on page 3-75. This window enables you to gain access to a FlowCollector by means of the NetFlow Collection Control option of the Tools menu.

  — If -i *is specified* in the PROGRAMFLAGS line, the FlowCollector bypasses (ignores) the authentication process.

In this case, you are connected to the FlowCollector, enabling you to begin using the NetFlow Collection Control window immediately without further requisite interaction (see the “Introduction to the NetFlow Collection Control Window” section on page 3-77).

- nf.resources file—Contains the variables and directory paths used to establish the FlowCollector startup and operating environment.

  The NFC_USERNAME variable in the nf.resources file establishes the password that the FlowCollector uses in authenticating a remote configuration/control request sent by the Analyzer to the FlowCollector (see the next section).

  The default value of the NFC_USERNAME parameter is “nfcuser.” This parameter can be set to any value, however. There is no default password.

User Name and Password Authentication Dialog Box

You can configure a FlowCollector either to enforce or to ignore the user name and password authentication requirement on startup, as described below.

If the FlowCollector has been configured to enforce user name and password authentication on startup, as described in the preceding section, the Display module prompts you to enter a user name and password in the User Name and Password Dialog box shown below.

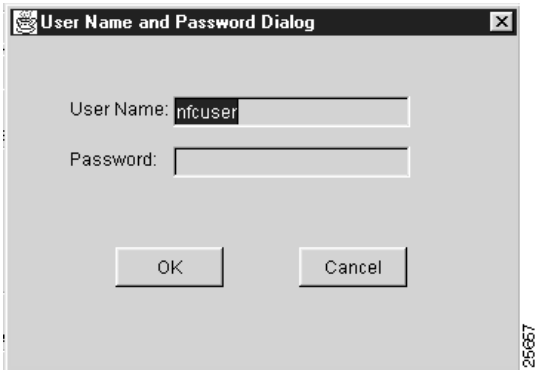

You must supply this information correctly and click OK.

During a connection attempt, the Analyzer sends the user name that you enter by means of this User Name and Password Dialog box to the FlowCollector for comparison against the user name specified in the NFC_USERNAME variable in the nf.resources file. If a match is found, the connection request to the FlowCollector is granted.

After the connection to the FlowCollector is completed, the Analyzer includes the user name and password in every configuration and control command sent to that FlowCollector.

The password that you enter is sent to the FlowCollector in clear text format. Upon receipt, the FlowCollector encrypts the password and compares it against the encrypted password for the user defined by the NFC_USERNAME variable in the /opt/CSCOnfc/config/nf.resources file of the FlowCollector. There is no default value for the password.

After you enter a valid user name and password for the FlowCollector, you can store the user name in the /opt/CSCOnfa/NFAUtility/config/NFCCC.txt file maintained by the UtilityServer. You save the user name by means of the User Validation window described in the “Saving the User Name” section on page 3-76. The FlowCollector host system password, however, is not stored in the NFCCC.txt file.

If the authentication process completes successfully, you are connected to the FlowCollector and the NetFlow Collection Control window appears immediately for use. You can then use this window to configure and control one or more FlowCollectors, as described in the "Introduction to the NetFlow Collection Control Window" section on page 3-77.

If you enter an invalid password in the User Name and Password Dialog box, the following error message appears in the User Validation window.

In this case, you must enter the correct password and click OK to proceed.

If you cannot open a connection to the FlowCollector because it is unreachable or not operational for some reason, the following Message window is displayed.

In this case, click OK and try again to connect to the FlowCollector. If you are unable to connect to the FlowCollector, determine the cause of the problem, if possible, and correct it.

Saving the User Name

A list of FlowCollectors is maintained by the UtilityServer in the NFCCC.txt file stored in its /opt/CSCOnfa/NFAUtility/config/ directory.

You establish the initial contents of this file when you install the UtilityServer. In so doing, you should take into account all of the FlowCollectors that you intend to control by means of the NetFlow Collection Control window.

The UtilityServer maintains the following parameters in the NFCCC.txt file for each FlowCollector:

- Collector_name (required)—The name of the FlowCollector.
- Port number (optional)—The port used by the FlowCollector in granting a connection to the Analyzer. If this parameter is not defined, 8550 is assumed by default.
- User_name (optional)—The user name for the FlowCollector. If this parameter is not defined, the User Name and Password Dialog window is displayed during a remote Analyzer connection attempt, enabling you to enter a valid user name for connecting to the FlowCollector.

If the user name and password that you enter for a FlowCollector passes the authentication check during an Analyzer connection request, you can store the user name in the UtilityServer's NFCCC.txt file by means of the User Validation window shown below.

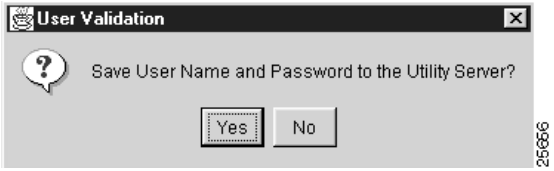

The capability to store the user name is provided as a user convenience in updating the NFCCC.txt file. If you forget the user name for the FlowCollector, you can view the contents of the NFCCC.txt file to refresh your memory.

For reasons of security, the host FlowCollector system password is not stored in the NFCCC.txt file. Nevertheless, the password that you enter to satisfy the FlowCollector authentication process remains in effect for the duration of the current Display module session. Thus, the FlowCollector "remembers" the user password until you exit from the current Display module session.

In the event that you exit the Display module and start again later, you must re-enter the appropriate password for any FlowCollector that is configured to enforce user authentication.

Introduction to the NetFlow Collection Control Window

If user authentication *is enabled* (the default state) on the FlowCollector to which you intend to connect, you will be prompted to specify a user name and password before being granted access to the FlowCollector.

When you complete the user authentication process, as described in the "User Name and Password Authentication Process" section on page 3-74, you are granted access to the NetFlow Collection Control window shown in Figure 3-53.

Figure 3-53 NetFlow Collection Control Window

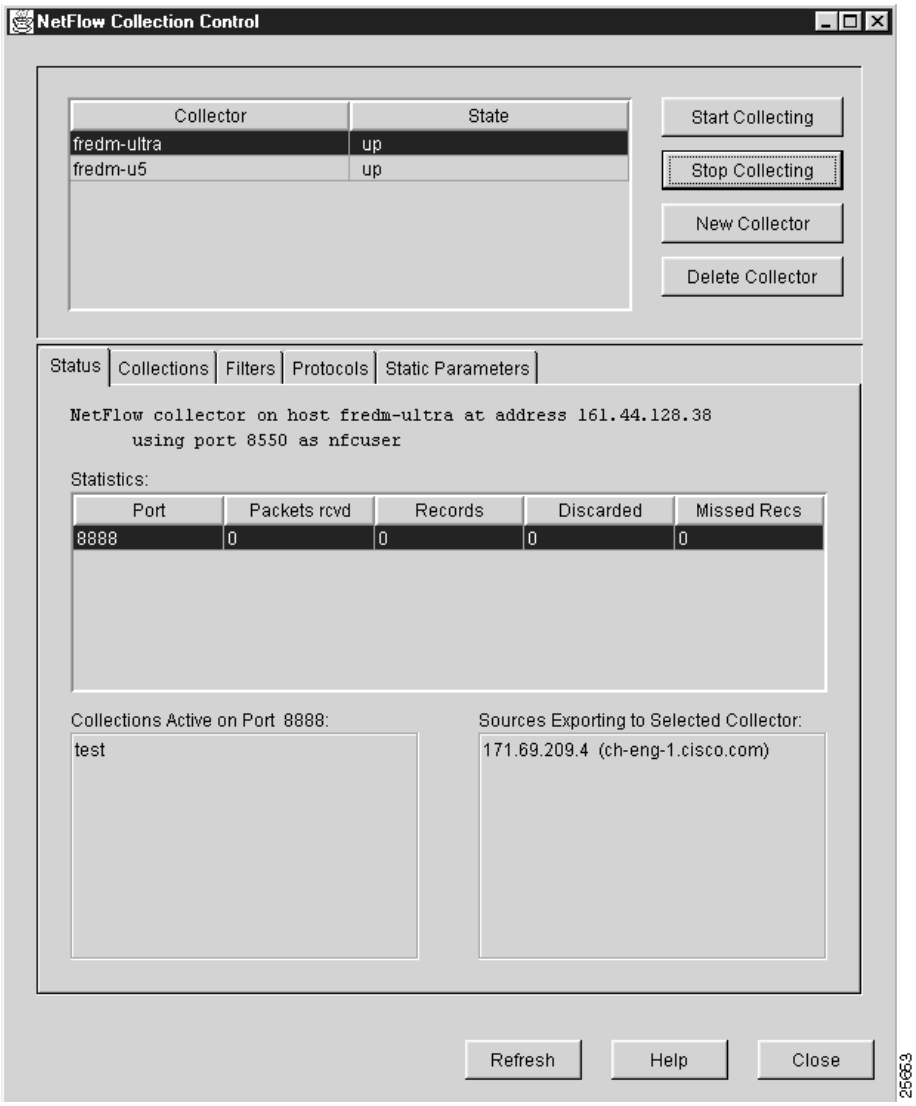

The rectangular area at the top of the NetFlow Collection Control window incorporates facilities that enable you to:

- Select a FlowCollector as the object of configuration and control commands

  FlowCollectors currently known to the UtilityServer are listed in a single-selection, “Collector” scroll list in the top left part of the window. You can select any member of this “Collector” list at any time as the object of FlowCollector configuration and control commands.

  The FlowCollectors included in the Collector list can vary to meet changing NetFlow data analysis requirements. The list can change from day-to-day, or even from one Display module session to another. The Analyzer enables you to add or delete FlowCollectors to meet your changing operational needs.

  The FlowCollectors displayed in this pane mirror those in the UtilityServer’s NFCCC.txt file. The NFCCC.txt file is maintained by the UtilityServer in its /opt/CSCOnfa/NFAUtility/config directory and is updated automatically when you add or delete a FlowCollector.

When you first access the NetFlow Collection Control window, the Collector list is populated with the names of the FlowCollectors currently defined in the Utility Server's NFCCC.txt file. By default, the first FlowCollector name in the NFCCC.txt file is automatically listed at the top of the Collector list and highlighted (selected). This FlowCollector remains in effect until you deliberately select another FlowCollector from the Collector list.

- Determine the state of each FlowCollector

  The state of each FlowCollector is indicated in the State list. The possible FlowCollector state flags are:

  — Up—The listed FlowCollector is accessible.

  — Down—The listed FlowCollector is not accessible.

  — Unknown—Indicates that no attempt has been made to establish a connection to the listed FlowCollector. In this case, click the name of the FlowCollector.

    When you select a FlowCollector in this state, a connection attempt is initiated. If user authentication is enabled on this FlowCollector, the UtilityServer attempts to get data from the FlowCollector. A pop-up labeled "User Name and Password For Collector....." appears. Enter the appropriate password in this dialog box and click OK.

    If you change the user name in this pop-up, a User Validation pop-up appears, querying you if you want to "Save User Name to Utility Server?" Click Yes. The user name is saved to the UtilityServer's NFCCC.txt file, and the corresponding FlowCollector flag in the State list is changed to "Up" or "Down," as appropriate

- Initiate or stop collections on FlowCollectors

  — Start Collecting button—Click Start Collecting to start the selected FlowCollectors. The Choose Collector window appears, enabling you to select and start any one or all of the FlowCollectors listed in the window.

After checking the appropriate boxes in the Choose Collector window, click OK to start the selected FlowCollectors.

Note This action initiates the flow of NetFlow data from the selected FlowCollectors to the UtilityServer.

Click Cancel to abandon the FlowCollector definition process and return to the NetFlow Collection Control window.

Click Clear to uncheck all of the boxes in the window, enabling you to check the boxes individually, as desired, to select only those FlowCollectors that you want to start. FlowCollectors so checked are started when you click OK.

— Stop Collecting button—Click Stop Collecting to stop the selected FlowCollectors.

The description above for the Start Collecting button applies in every respect for the Stop Collecting button, except that the latter stops the selected FlowCollectors.

— New Collector button—Click the New Collector button to add a new FlowCollector to the Collector list.

The New Collector Dialog pop-up appears, enabling you to enter the name and port address of the new FlowCollector.

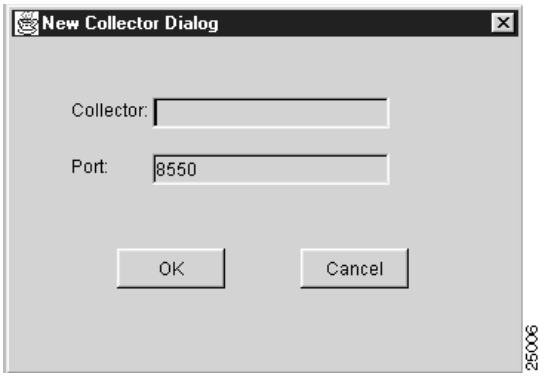

When you click OK, the new FlowCollector is added to the Collector list (see Figure 3-54).

Click Cancel to abandon the new FlowCollector definition process and return to the NetFlow Collection Control window.

Figure 3-54 Result of Adding a New FlowCollector

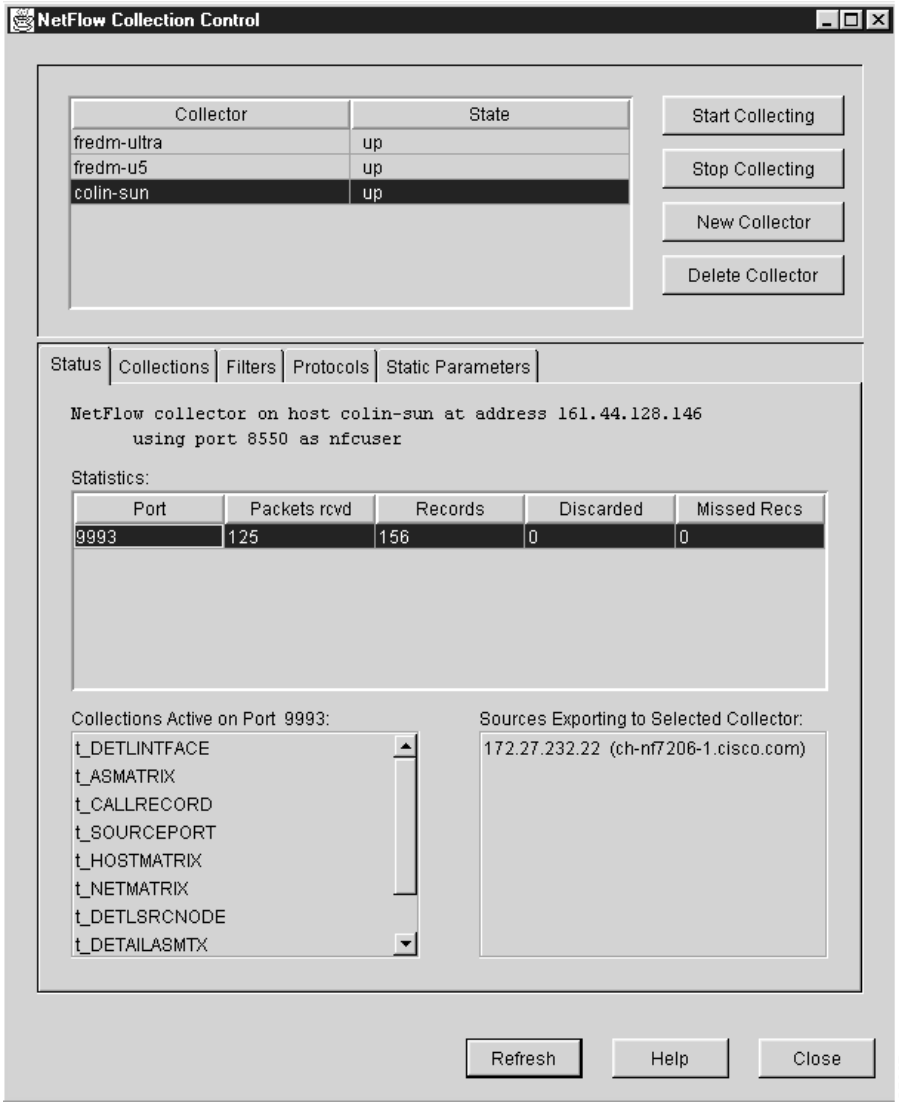

— Delete Collector button—Click Delete Collector to remove the currently selected FlowCollector from service in your network.

Suppose, for example, that you want to delete the FlowCollector named fredm-ultra from the Collector list. Make sure that fredm-ultra is currently selected in the Collector list. Click Delete Collector to access the following window, which asks you for confirmation of your intent to delete the selected FlowCollector.

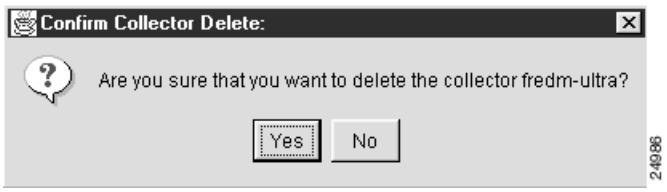

Click “Yes” to delete the selected FlowCollector and return to the NetFlow Collection Control window. Click “No” to abandon the operation and return to the NetFlow Collection Control window.

The rectangular area at the bottom of the NetFlow Collection Control window contains several tabs, any one of which you can select at any time to perform desired NetFlow collection control tasks for a selected FlowCollector. The panels associated with these tabs are described in the sections referenced below:

- Status panel—See the “Using the Status Panel” section on page 3-84.
- Collections panel—See the “Using the Collections Panel” section on page 3-86.
- Filters panel—See the “Using the Filters Panel” section on page 3-92.
- Protocols panel—See the “Using the Protocols Panel” section on page 3-98.
- Static Parameters panel—See the “Using the Static Parameters Panel” section on page 3-102.

In addition, the bottom of the NetFlow Collection Control window contains the following selectable push-buttons:

- Refresh button—Click Refresh to update all the information currently being displayed in the NetFlow Collection Control window for the selected FlowCollector.
- Help button—Click Help to activate an embedded Help system in the Analyzer that assists you in performing a variety of Analyzer configuration, control, and display tasks.
- Close button—Click Close to close the NetFlow Collection Control window and return to the main Display module window.

Globally Replicating FlowCollector Configuration Parameters

You can replicate the configuration and control parameters that you define for one FlowCollector into other FlowCollectors by means of the Choose Collectors window shown in Figure 3-55.

Figure 3-55 Choose Collectors Window

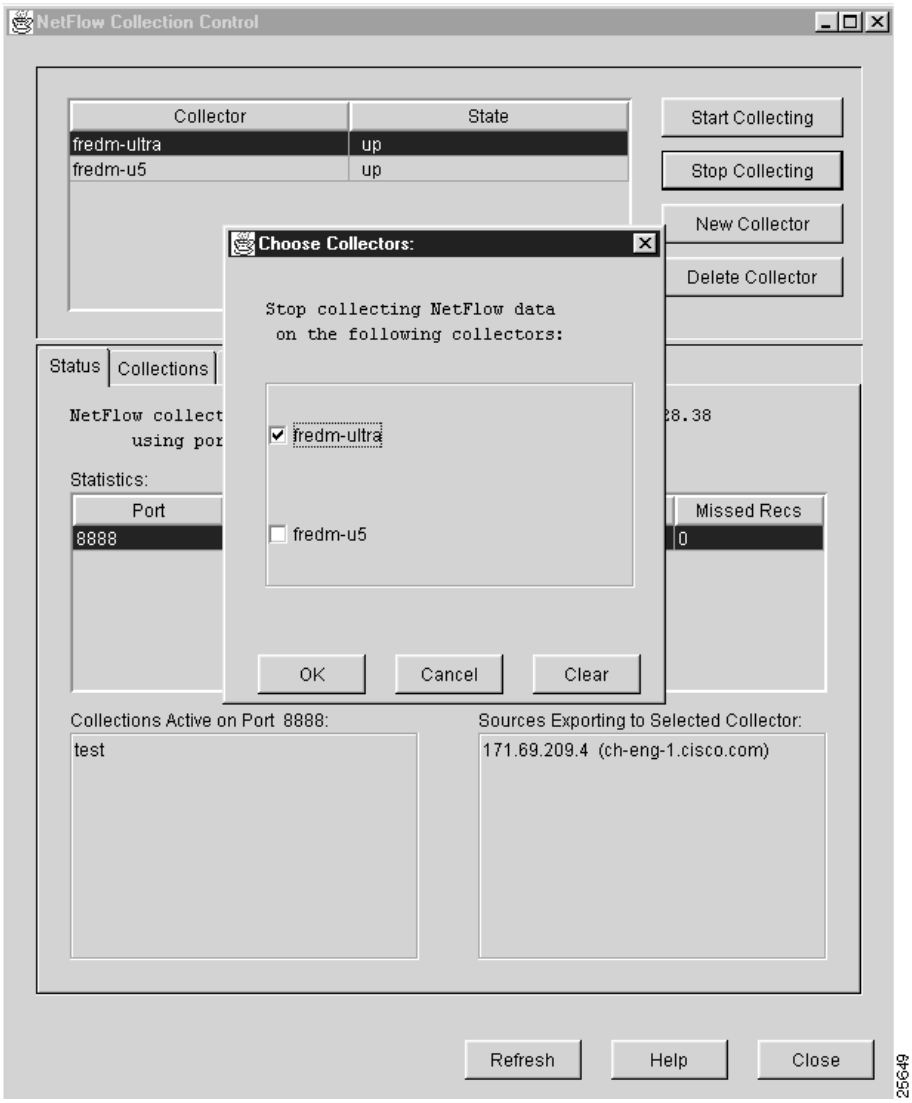

The Choose Collectors window appears automatically each time you attempt to perform an operation that can be logically applied to other FlowCollectors, such as when you:

- Start a FlowCollector
- Stop a FlowCollector
- Create collections, filters, and protocols
- Delete collections, filters, and protocols
- Activate or deactivate a collection of the same name on multiple FlowCollectors

From Figure 3-55, note that:

- The FlowCollectors listed in the Choose Collectors window are the same as those listed in the Collector list of the NetFlow Collection Control window.
- Each FlowCollector listed in the Choose Collectors window has an associated checkbox.

When you click a checkbox for one or more of the FlowCollectors listed in the Choose Collectors window, the associated FlowCollector is selected as a target for a NetFlow configuration and control task.

When you click OK in the Choose Collectors window, the intended operation for the selected FlowCollectors is carried out on a “global” scale. The Choose Collectors window is then closed and you are returned to the NetFlow Collection Control window.

If you attempt to perform a configuration or control task for a selected FlowCollector for which the intended operation is irrelevant (such as attempting to start a collection process on a FlowCollector that is already running), an appropriate error message appears, such as that shown below.

If you click Cancel in the Choose Collectors window, the selection process is abandoned, closing the Choose Collectors window, and returning you to the NetFlow Collection Control window.

If you click Clear in the Choose Collectors window, all the boxes in the window are unchecked, enabling you to reconsider your selections.

Using the Status Panel

A typical Status panel in the NetFlow Collection Control window is shown in Figure 3-56.

Figure 3-56 Sample Status Panel of NetFlow Collection Control Window

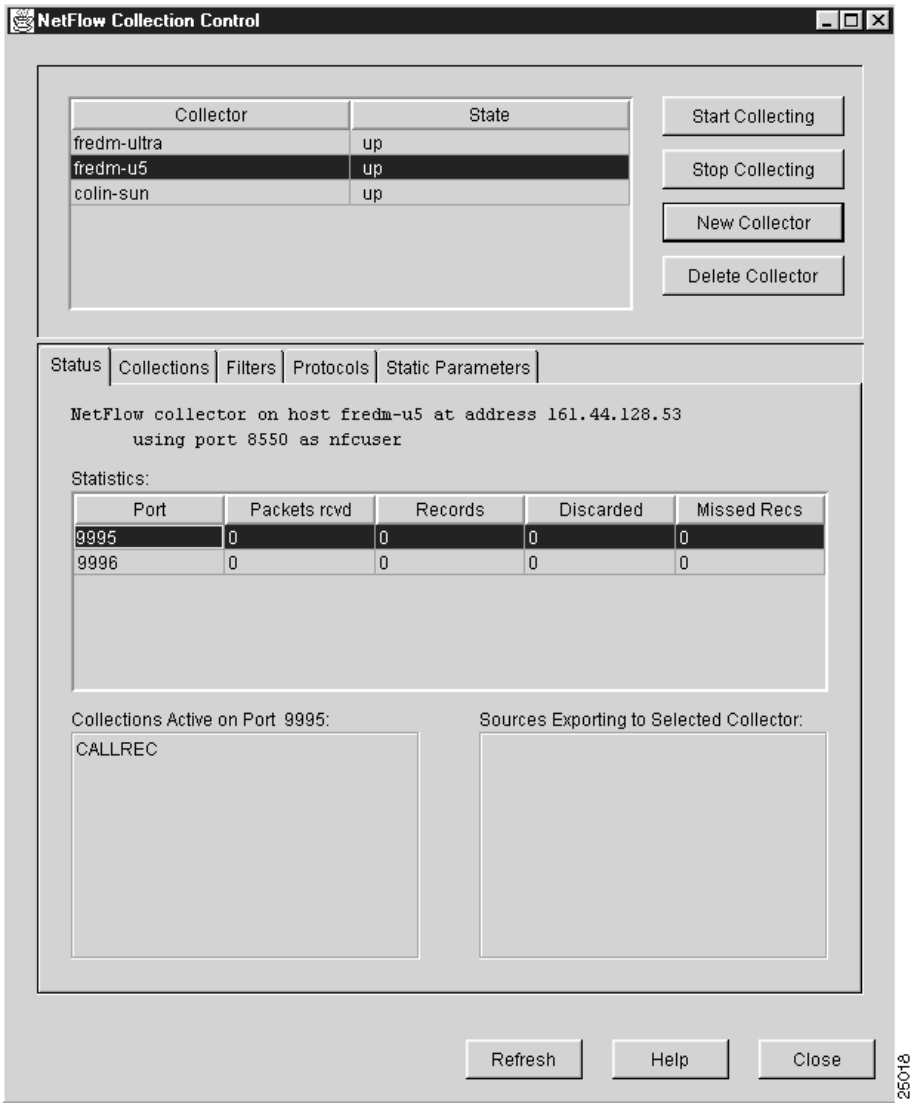

The Status panel presents information about active collections on the selected FlowCollector. The panel includes the following facilities:

- Status text—Provides the following information about the selected FlowCollector:
  - Name of the workstation on which the FlowCollector is operating
  - IP address of the FlowCollector
  - Application session port number being used by the FlowCollector
  - FlowCollector user name
- Statistics pane—Provides a snapshot of current per-port traffic statistics for the selected FlowCollector. These statistics include the following:
  - Port—Displays the active UDP ports on which the FlowCollector is listening for UDP datagrams
  - Packets rcvd—Displays the number of packets received during the flow

  - Records—Displays the number of records received during the flow
  - Discarded—Displays the number of records discarded during the flow
  - Missed Recs—Displays the number of records missed during the flow
- Collections Active on Port XXXX pane—Lists the collection processes that are active on the port number (XXXX) selected in the Statistics pane.

  To view the collections active on a port other than the currently selected one, select the desired port.
- Sources Exporting to Selected Collector pane—Lists the network devices currently exporting NetFlow data to the selected FlowCollector.

  This pane supports no user interaction. Its purpose is to provide IP address and logical device name information for all the devices currently exporting NetFlow data to the selected FlowCollector.

  Each row in this pane contains the IP address and the names of the NetFlow exporting devices (enclosed within parentheses opposite the IP address).

  Multiple devices can export NetFlow data to any given FlowCollector. Therefore, there are always as many rows in this table as there are network devices exporting NetFlow data to the selected FlowCollector.

Note Once you configure the selected FlowCollector to operate using the facilities provided by the NetFlow Collection Control windows, you ordinarily keep the Status panel selected until you need to use other panels to perform various FlowCollector configuration, control, and display tasks.

Using the Collections Panel

The Collections panel lists the collections currently defined for the selected FlowCollector. As shown in Figure 3-57, this panel enables you to accomplish the following tasks:

- Define new collections for a selected FlowCollector
- Delete existing collections from a selected FlowCollector
- Modify the parameters of existing collections for a selected FlowCollector

Each time you change the configuration parameters for a selected FlowCollector, a command is sent to that FlowCollector to update its nfconfig.file, thus causing it to operate according to the new or changed configuration parameters.

One of the key parameters that you can set in this window is the collection filter. A filter can incorporate multiple permit and deny attributes (such as source addresses, interfaces, and so on), but all such attributes must be related to one filter name. You can apply a filter with multiple permit and deny attributes to multiple collections, but any given collection can have at most just one filter associated with it at any time.

Figure 3-57 Sample Collections Panel of NetFlow Collection Control Window

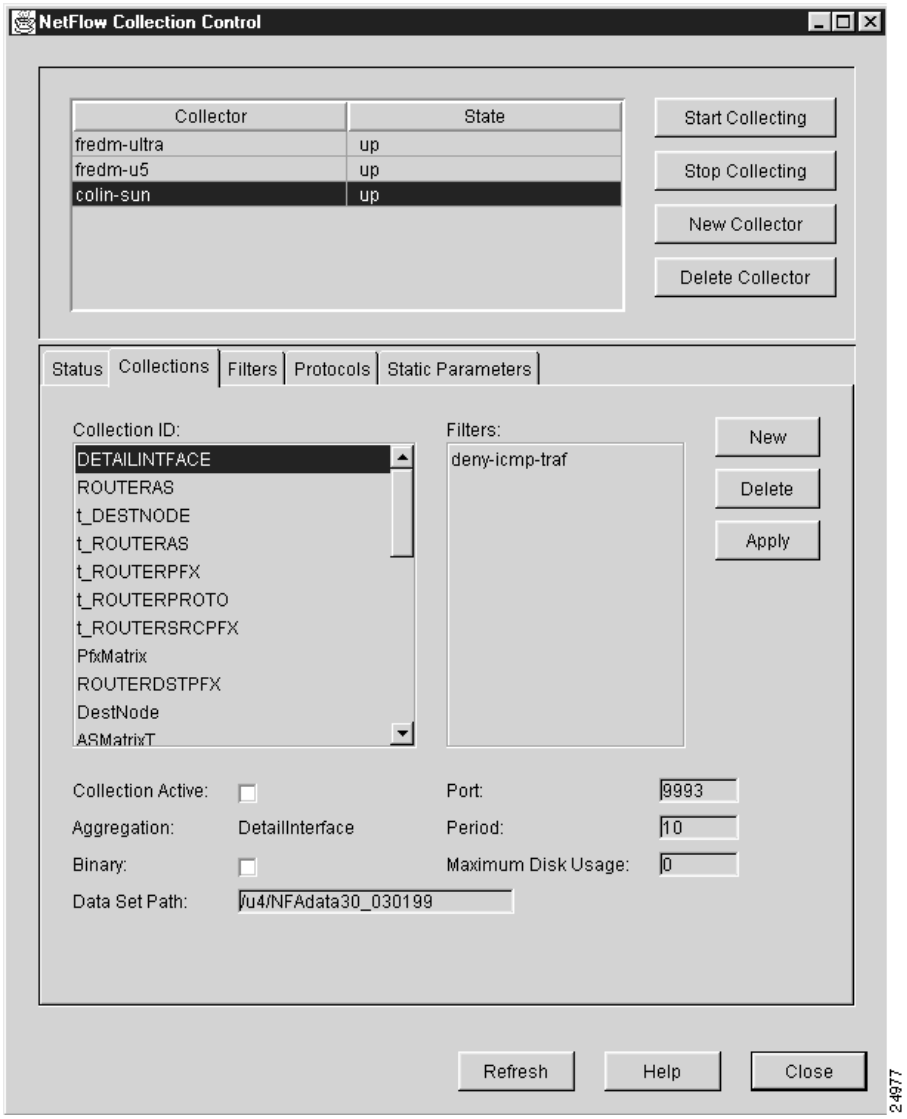

The Collections panel provides the following FlowCollector configuration facilities:

- New button—Click this button to access a secondary New Collection pane that enables you to define a new collection. For more information the New Collection pane, refer to the "Using the New Collections Panel" section on page 3-89.
- Delete button—Click this button to delete the currently selected collection.

  To delete a collection, select the name in the Collection ID pane of the collection that you want to delete. Then, click Delete. A command is then sent to the selected FlowCollector to delete this collection from the FlowCollector nfconfig.file.
- Apply button—Click this button to put into effect any changes that you have made in the configuration parameters for the selected collection.

  To modify the parameters for the selected collection, edit the appropriate text fields and check boxes in the Collections panel, as desired.

When you click Apply, a command is sent to the selected FlowCollector to update the configuration parameters in its nfconfig.file, altering the behavior of the FlowCollector according to the configuration changes made in the Collections panel.

The FlowCollector stores the collected data and restarts the collection period when the command is processed.

- Collection ID pane—This single-selection scroll pane provides a mechanism for selecting a collection name as the object of the following tasks:
  - Viewing collection parameters
  - Altering collection parameters
  - Deleting a collection

You can select only one collection name at a time in this pane.

- Filters pane—This single-selection scroll pane lists the filters currently available for use with the collections listed in the Collection ID pane.

  If the Filters scroll pane contains no entries and you want to use a filter with a collection, you must define one or more filters for use with selected FlowCollectors and collections. Such filters are defined by means of the Filters panel, which is described in the "Using the Filters Panel" section on page 3-92.

- Collection Active check box—Use this check box to activate or deactivate a collection. Checking the box immediately activates the collection name selected in the Collection ID scroll pane; unchecking the box deactivates the selected collection name.

  You can have as many as 10 collections active for any given Flowcollector.

  You can use the Collection Active check box in conjunction with the Choose Collector window shown in Figure 3-55 to activate or deactivate collections of the same name on two or more FlowCollectors.

  For example, assume that there are three FlowCollectors in your network named fredm-ultra, fredm-u5, and colin-sun, as shown in Figure 3-57, and that the collection named "DetailIntface" is selected in the Collection ID scroll pane for the FlowCollector named colin-sun.

  Assume further that a collection named "DetailIntface" is also defined for the FlowCollectors named fredm-ultra and fredm-u5. (Note that the attributes of individual collections for different FlowCollectors may differ, even though the collections have the same name.)

  If you checked the Collection Active check box for the selected FlowCollector named colin-sun and clicked Apply, the Choose Collector window would appear, enabling you to select (check) any one or all of the other FlowCollectors listed in the Choose Collector window (each of which has a collection named "DetailIntface") and to activate that named collection on all selected FlowCollectors.

  This "global" effect of the Collection Active check box enables you to conveniently activate and deactivate multiple collections of the same name across multiple FlowCollectors in your network.

  This global activation capability works only in conjunction with the Collection Active check box. If you change any parameter in the Collections panel other than the state of the Collection Active check box before you click Apply, the Choose Collector window will not appear, and you will not be able to activate or deactivate multiple collections simultaneously.

- Aggregation name—This field identifies the name of the aggregation scheme associated with the selected collection.

- Binary check box—Use this selection box to determine whether or not the collected data for the selected collection is to be stored in binary form.

  Checking this box causes the data files for the selected collection to be stored by the selected FlowCollector in binary form; unchecking the box causes the data files to be stored in ASCII form. The default setting for this box is unchecked.

  Choosing the binary file format reduces network traffic data

- Data Set Path field—Use this text field to establish the directory path to the selected FlowCollector in which data files for the new collection are to be stored. This path must be writable by the FlowCollector.
- Port field—Use this text field to identify the UDP port number (in the range from 1024 to 65535) used by the selected FlowCollector to listen for UDP export datagrams. You can have as many as ten collections active on a given port.
- Period field—Use this text field to establish the period or granularity (in minutes) between collections for the selected collection. The FlowCollector creates one aggregated data collection file for each period. The value in this field can range from 1, the default value, to 1440 (24 hours).
- Maximum Disk Usage field—Use this text field to indicate the number of bytes of disk space allowable (in MB) before the selected FlowCollector starts to delete old data files to make room for new data files. The default value of this field is 0.

Using the New Collections Panel

When you click the New button in the Collections panel, a New Collection panel (see Figure 3-58) appears in place of the Collections panel.

Figure 3-58 Sample New Collection Panel of NetFlow Collection Control Window

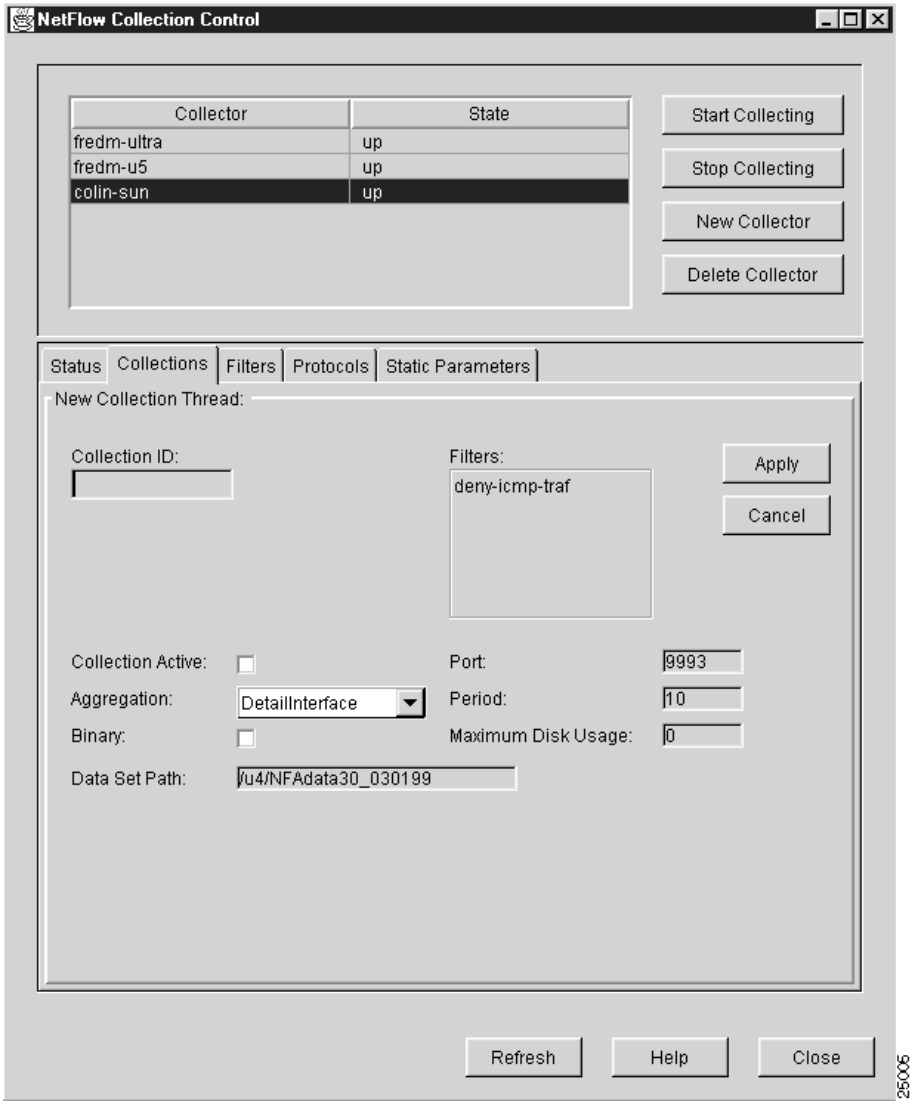

When you click New in the Collections panel, all of the parameters of the most recently selected Collections panel are propagated into the New Collections panel, giving you a basis for:

- Accepting the existing collection parameters and associating them with a new collection ID.
- Changing the existing collection parameters and associating them with a new collection ID name.

In either case, the New Collections panel:

- Incorporates all of the collection configuration facilities available to you in the Collections panel.
- Incorporates a Collection ID pane, enabling you to define a unique name for the new collection process.

When you specify a new collection for the selected FlowCollector, click OK to put the new collection into effect. Doing so sends configuration commands to the selected FlowCollector, updating its nfconfig.file and causing it to recognize the new collection process in aggregating NetFlow data.

If you decide to abandon the new collection definition process, click Cancel to return to the Collections panel.

The New Collection panel provides the following configuration facilities for defining new collections:

- Collection ID—Use this text field to define a unique alphanumeric name for identifying the new collection. You can specify up to 14 characters for the collection name, with no intervening spaces.
- Filters—Use this single-selection scroll pane to select a filter that you want to apply to the new collection.

  You can choose only one of the filters listed in this scroll pane for use with the new collection.

  To define filters by means of the Filters panel, refer to the “Using the Filters Panel” section on page 3-92.
- Collection Active—Use this check box to set the new collection to either an active or an inactive state.

  To set the new collection to an active state, check this box; to set the new collection to an inactive state, leave this box unchecked.
- Aggregation—Use this single-selection scroll pane to select the desired aggregation scheme for the new collection.

  Click the down arrow in this pane to access a scroll list, which enables you to select the desired aggregation scheme for the new collection.
- Binary—Use this check box to determine whether or not you store the data files for the new collection in binary form.

  Check this box to cause the data files for the new collection process to be stored in binary form; leave the box unchecked to store the data files for the new collection in ASCII form (the default).
- Data Set Path—Use this text field to establish the directory path to the selected FlowCollector in which data files for the new collection are to be stored. This path must be writable by the FlowCollector.
- Port—Use this text field to define the UDP port number (in the range from 1024 to 65535) on which the selected FlowCollector listens for UDP export datagrams generated by the new collection process.
- Period—Use this text field to establish the period (in minutes) between collections for the new collection task. The value in this field can range from 1, the default value, to 1440 (24 hours).
- Maximum Disk Usage—Use this text field to indicate the number of bytes of disk space allowable (in MB) before the selected FlowCollector starts to delete existing data files for the current collection to make room for new data files. The default value of this field is 0.

The New Collection panel also includes the following collection control facilities:

- Apply button—Click this button to put into effect the configuration information that you have established in the various fields and check boxes of the New Collection panel.

  When you click Apply, a command to the selected FlowCollector to establish the appropriate configuration parameters in its nfconfig.file for the new collection.
- Cancel button—Click Cancel at any time to abandon the new collection creation task and return to the Collections panel.

Using the Filters Panel

When you click the Filters tab of the NetFlow Collection Control window, the Filters panel appears (see Figure 3-59).

Figure 3-59 Sample Filters Panel of NetFlow Collection Control Window

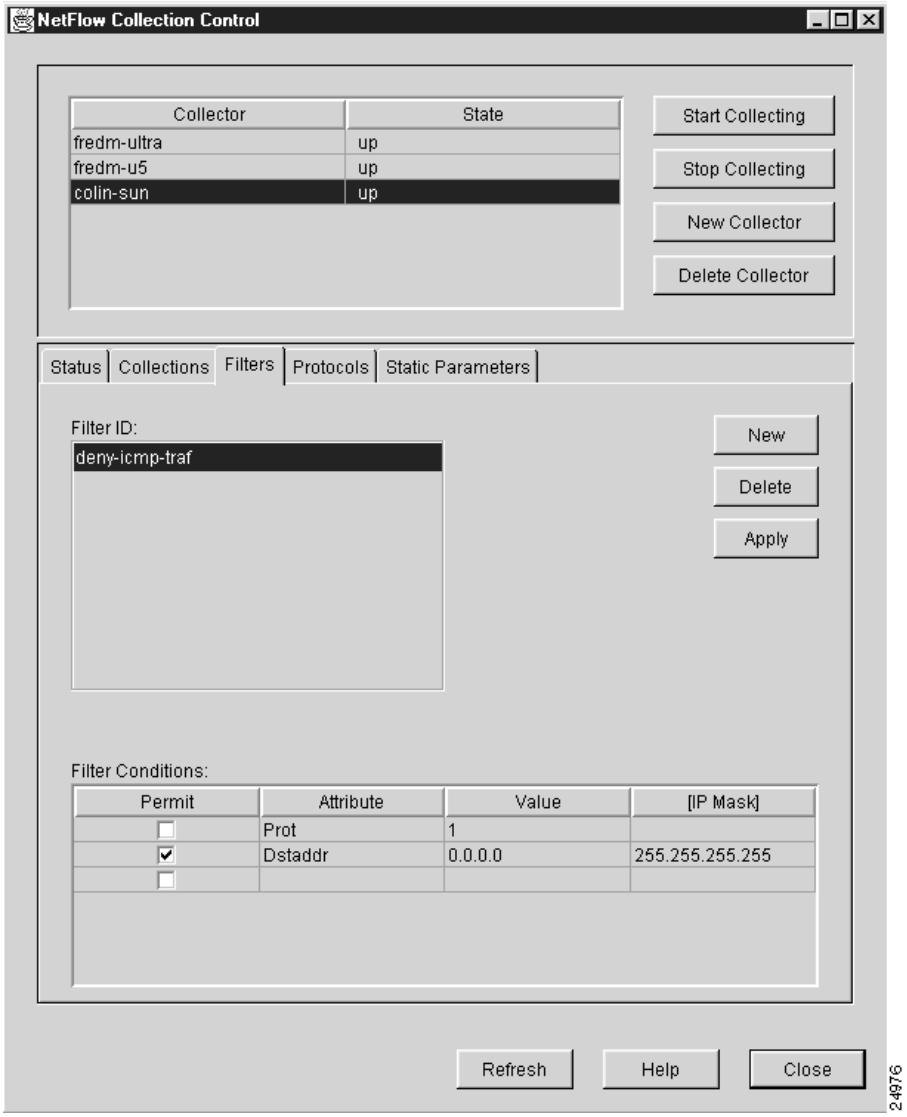

The Filters panel lists the filters currently defined for the selected FlowCollector and incorporates facilities that enable you to:

- Create a new filter
- Delete an existing filter
- Change existing filter parameters

When you create a new filter, modify an existing filter, or delete an existing filter, a command is sent to the selected FlowCollector, causing it to update its nfconfig.file and to operate accordingly.

Note You cannot modify a filter that is currently being used by a FlowCollector. You must remove the filter from any collections for which it is defined and then make desired modifications.

The Filters panel provides the following filter display and definition facilities:

- Filter ID scroll pane—This single-selection scroll pane lists the filters currently defined for the selected FlowCollector.

  You use filters to permit or deny the collection of certain NetFlow data for the selected FlowCollector, based on specified traffic attributes.

  When you select the name of a filter in the Filter ID pane, the configuration parameters defined for that filter are displayed in the Filter Conditions pane. Each time you select a different filter, the information in the Filter Conditions pane is updated.

- Filter Conditions pane—This pane consists of a table that displays the various filter attributes and associated parameters that are defined for the currently selected filter.

  As shown in Figure 3-59, you can define multiple attributes for any given filter, enabling you to tailor the collection of NetFlow traffic data according to specified attributes.

  The filter attributes that you can permit or deny for a given filter are listed in Table 3-5.

Table 3-5 Filter Attributes and Associated Values

Filter Attribute	Filter Value
Srcaddr Dstaddr Nexthop	The IP addresses, respectively, of the traffic source, the traffic destination, and the next hop device routing the traffic to the destination. Each IP address requires a network mask.
Srcport Dstport	The assigned port number for the transport layer protocol (RFC 1700) at the traffic source and destination, respectively. The port number can range from 1 to 65535.
Srcinterface Dstinterface	The numeric identifier of the physical interface at the source and destination, respectively.
Prot	The protocol number in the flow record, as specified in the /etc/protocols file. The protocol number can range from 1 to 255.
Protocol	The protocol name, as displayed in the Protocol pane of the Protocol panel (see Figure 3-62) and defined in the nfknown.protocols file.
TOS	The type-of-service (ToS) byte (which includes IP precedence and Type of Service fields) provides a way to prioritize traffic. The value of the ToS byte can range from 0 to 255.
SrcAS DstAS	The source and destination autonomous system (AS) number, respectively.

The Filter Conditions pane incorporates the following filter definition facilities:

— Permit column—Includes a check box in each row, enabling you to alter the sense of the selected filter, that is, whether to permit or deny the collection of NetFlow data, based on the currently defined filter attributes.

If you check the Permit box, the collection of NetFlow data according to the associated filter attribute is permitted. If you uncheck the Permit box, the collection of data according to the associated filter attribute is denied.

— Attribute column—Lists the specific filter attributes that have been defined for the selected filter.

As noted above, multiple permit/deny attributes can be specified for a selected filter. Typically, you assign a desired subset of the filter attributes listed in Table 3-5 to a given filter, thereby tailoring FlowCollector operations to suit your data analysis requirements.

— Value column—Contains a value (see Table 3-5) that depends on the specified filter attribute.

If you specify Srcaddr, Dstaddr, or Nexthop as a filter attribute, you would specify a valid IP address in the associated value field of the row, since all three filter attributes are IP-address based parameters.

You can specify any valid IP address in the value field. Similarly, if you specify Srcport as a filter attribute, the value field in that row would contain the source port number.

Any filter attribute that requires an IP address in the value field can also accept an IP subnet mask in the associated [IP Mask] column (see below).

Conversely, if you specify Srcport as a filter attribute, the [IP Mask] field is ignored, since the port number is not an IP-based attribute.

— [IP Mask] column—Contains an IP mask for the associated value field, if appropriate.

You set the [IP Mask] field only if the associated filter attribute is an IP-address based parameter, such as Srcaddr, Dstaddr, or Nexthop.

The subnet mask that you specify in the [IP Mask] field depends on how you choose to define subnetting in your network, that is, how many bits in the IP address you want to be in effect for NetFlow data filtering purposes.

For example, a subnet mask of 255.255.255.0 treats the first 12 bits of an IP address as relevant for NetFlow data filtering purposes. In specifying the subnet mask, however, the value that you enter in the [IP Mask] field is the 1's complement of the subnet mask. Thus, for a subnet mask of 255.255.255.0, you would specify 0.0.0.255 in this field.

- New button—When you click New, a New Filter pane replaces the current Filters pane, enabling you to create a new filter.

  As time passes, it may be desirable to define new filters to serve new NetFlow data analysis requirements. Hence, the ability to define new filters is provided, as described in the "Using the New Filters Panel" section on page 3-95.

- Delete button—When you click Delete, the currently selected filter is deleted from the currently selected FlowCollector.

  At the same time, a command is sent to the affected FlowCollector, instructing it to update its nfconfig.file to reflect the deletion of the selected filter.

- Apply button—When you click Apply, any changes that you have made in the fields of the Filter Conditions pane are put into effect, creating a new set of parameters for the selected filter.

  You can edit the text fields of the Filter Conditions pane, check or uncheck Permit boxes, or select new filter attributes for assignment to the selected filter.

  When you have changed all areas of the Filter Conditions pane as desired, click Apply to recreate the filter with the new configuration parameters.

Using the New Filters Panel

When you click the New Filter button in the Filters panel, a New Filter panel (see Figure 3-60) appears in place of the existing Filters panel.

Figure 3-60 Sample New Filters Panel of NetFlow Collection Control Window

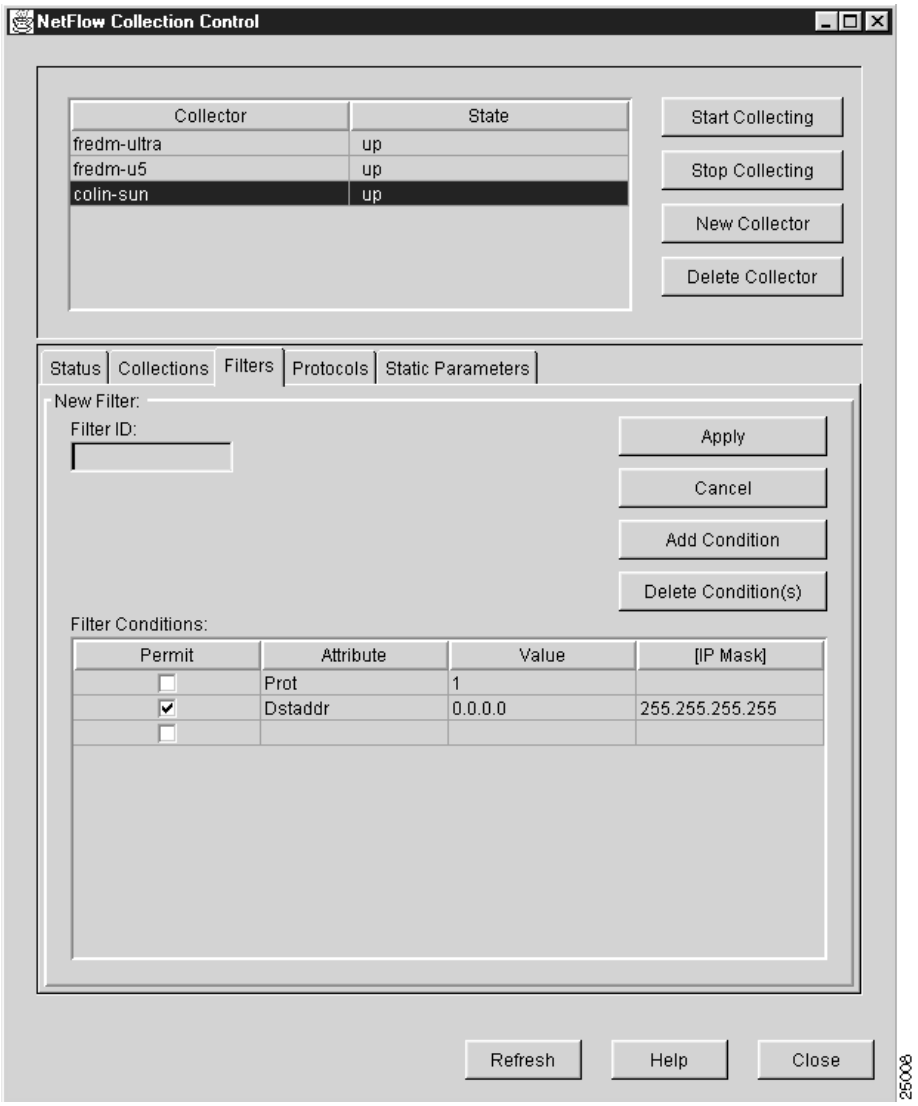

All of the parameters of the most recently selected Filters panel are propagated into the New Filter panel, giving you a basis for:

- Accepting the existing filter parameters and associating them with a new filter name.
- Changing the existing filter parameters and associating them with a new filter name.

In either case, the New Filter panel incorporates all of the filter definition facilities available to you through the Filters panel. In addition, the New Filter panel incorporates a Filter ID pane that enables you to define a unique name for the new filter.

When you create a new filter for the selected FlowCollector, a command is sent to the FlowCollector to update its nfconfig.file accordingly.

The New Filter panel provides the following new filter definition facilities:

- Filter ID text field—This text field provides a facility for creating a unique name for a new filter.

  You can specify a new filter name consisting of up to 14 alphanumeric characters with no intervening spaces.

- Filter Conditions pane—This pane contains a table that enables you to define the various attributes for a new filter.

  If you want to define multiple filter attributes for a new filter using this table, you can click the Add Condition button to add rows to the table.

  The Filter Conditions pane incorporates the following filter configuration facilities:

  — Permit column—Includes a check box in each row that enables you to permit or deny the associated filter attribute.

    If you check the Permit box, the collection of NetFlow data according to the associated filter attribute is permitted. If you uncheck the Permit box, the collection of such data is denied.

  — Attribute column—Enables you to select the filter attributes that you want to define for the new filter.

    You can define multiple filter attributes for a new filter. Typically, you assign to the new filter a desired subset of the filter attributes listed in Table 3-5. Thus, you can tailor the selected FlowCollector to operate in a way that meets your NetFlow data analysis requirements.

  — Value column—Contains a value (see Table 3-5) that depends on the specified filter attribute.

    For example, if you specified Srcaddr, Dstaddr, or Nexthop as a filter attribute, you would also specify a valid IP address in the associated value field for that row, since all three filter attributes are IP-address based parameters. You can specify any valid IP address in the value field.

    Similarly, if you specified Srcport as a filter attribute, you would also specify the number of the source port in the value field for that row.

    A filter attribute that calls for an associated IP address in the value field, such as Srcaddr, Dstaddr, or Nexthop, can also accept an IP subnet mask in the associated [IP Mask] field of that row (see the description of the [IP Mask] column below).

    Conversely, if you specified Srcport as a filter attribute, the [IP Mask] field would be ignored, since the port number is not an IP-based filter attribute.

  — [IP Mask] column—Contains an IP mask for the associated value field, if appropriate.

    You set the [IP Mask] field only if the associated filter attribute is an IP-address based parameter, such as Srcaddr, Dstaddr, or Nexthop.

    The subnet mask that you specify in the [IP Mask] field depends on how you choose to define subnetting in your network (the number of bits in the IP address that you want to put into effect for NetFlow data filtering purposes).

    For example, a subnet mask of 255.255.255.0 treats the first 12 bits of an IP address as relevant for NetFlow data filtering purposes. In specifying the subnet mask, however, what you enter in the [IP Mask] field is the 1's complement of the subnet mask. Thus, for the subnet mask 255.255.255.0, you would enter 0.0.0.255 in this field.

- Apply button—When you have entered all of the new filter configuration parameters, including a unique Filter ID, click Apply to put the new filter attributes into effect.

  When you click Apply, a collection creation command is sent to the selected FlowCollector, updating its nfconfig.file and causing it to recognize the new filter.

- Cancel button—Click Cancel to abandon the new filter creation task and return to the Filters panel.
- Add Condition button—To define additional filters, click Add Condition as many times as needed to create the desired number of rows in the Filter Conditions table.

  Each activation of the Add Condition button creates an additional row in the table for filter definition. Figure 3-61, for example, shows the addition of three rows for defining filter attributes.

Figure 3-61 Add Condition Function of New Filter Window

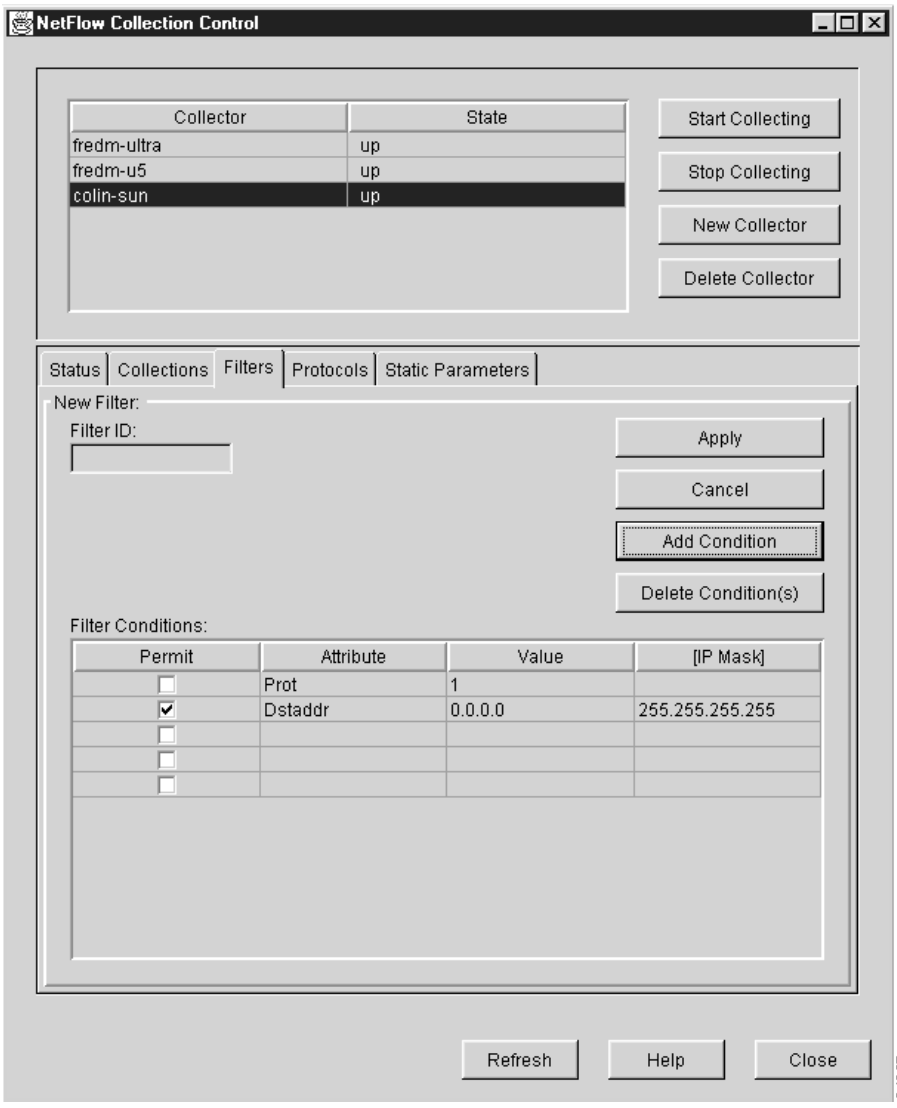

- Delete Conditions(s) button—To delete a filter definition from the Filter Conditions table, select the row in the table that you want to delete and click Delete Condition(s).

  Repeat this process as many times as necessary to delete the desired number of filter definitions from the Filter Conditions table.

Using the Protocols Panel

When you click the Protocols tab of the NetFlow Collection Control window, the Protocols panel appears (see Figure 3-62).

Figure 3-62 Sample Protocols Panel of NetFlow Collection Control Window

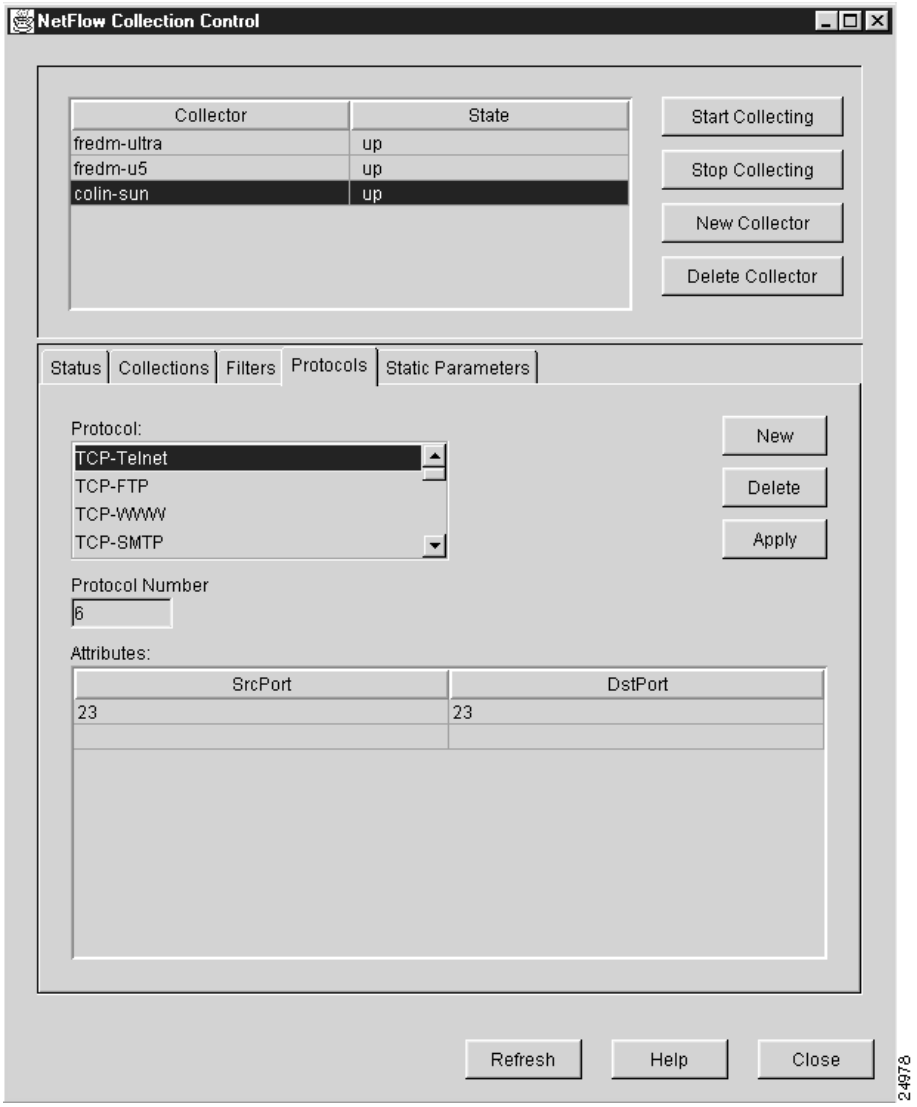

The Protocols panel lists the protocols currently defined for the selected FlowCollector, enabling you to:

- Create a new protocol
- Delete an existing protocol
- Change the parameters of an existing protocol

When you create a new protocol, modify an existing protocol, or delete an existing protocol, a command is sent to the selected FlowCollector to update its nfknown.protocols file with the current protocol definitions and port numbers.

The Protocol panel provides the following protocol display and definition facilities:

- Protocol pane—This single-selection scroll pane lists the protocols currently defined for the selected FlowCollector.
- Protocol Number text field—This field contains the protocol number of the selected protocol.

  Different protocols can use different ports while sharing the same protocol number. For example, FTP is protocol 6 and uses ports 20 and 21; similarly, Telnet is also protocol 6, but uses port 23.

  The protocol numbers and port assignments appearing in the Protocol Number field and Attributes pane, respectively, reflect standard assignments regulated by the Internet Assigned Numbers Authority (IANA). It is strongly recommended that you adhere to these standard definitions to ensure proper Analyzer operation.
- Attributes pane—Contains the source port number and the destination port number of the devices that are communicating by means of the selected protocol. Network devices communicate with each other by means of well-known protocols and well-known ports.

  The port numbers reflected in the Attributes pane influence the way the FlowCollector translates and stores NetFlow data files. For example, if the selected protocol is protocol “6” (for TCP) in the Protocol Number field and the SrcPort and the DestPort are using port “20” or “21” (for FTP), the stored data files will reflect the translated notation “TCP-FTP” when the data files are written by the FlowCollector for the associated traffic flow. In this manner, the port numbers appearing in the Attributes pane are mapped to the selected protocol name.

  As you can see from the protocols listed in the Protocols pane of Figure 3-62, different network protocols can have the same protocol number, but such protocols can be associated with different ports, depending on the nature of the communicating applications.

  In general, port numbers are associated with application-level protocols; you should exercise extreme care in modifying any existing pre-determined port numbers.
- New button—When you click New, a New Protocol pane replaces the current Protocol pane, enabling you to create a new protocol definition.

  In certain operating circumstances, it may be desirable to define a new protocol to run on certain ports to serve some special or proprietary need. Accordingly, the New Protocol panel provides the ability to define new protocols and associated port numbers. This facility is described in the “Using the New Protocol Panel” section on page 3-100.
- Delete button—When you click Delete, the currently-selected protocol definition is deleted from the currently selected FlowCollector.

  Simultaneously, a command is sent to the affected FlowCollector, causing it to update its nfknown.protocols file to reflect the deletion of the selected protocol.
- Apply button—When you click Apply, any changes that you have made in the fields of the Protocols panel are put into effect, thereby creating a new protocol definition for the selected FlowCollector.

  You can edit the configuration parameters in the Protocols panel. When you complete the desired changes in this panel, click Apply to recreate the protocol definition with the new configuration parameters.

Using the New Protocol Panel

When you click the New button in the Protocols panel, a New Protocol panel (see Figure 3-63) appears in place of the existing Protocols panel.

Figure 3-63 Sample New Protocol Panel of NetFlow Collection Control Window

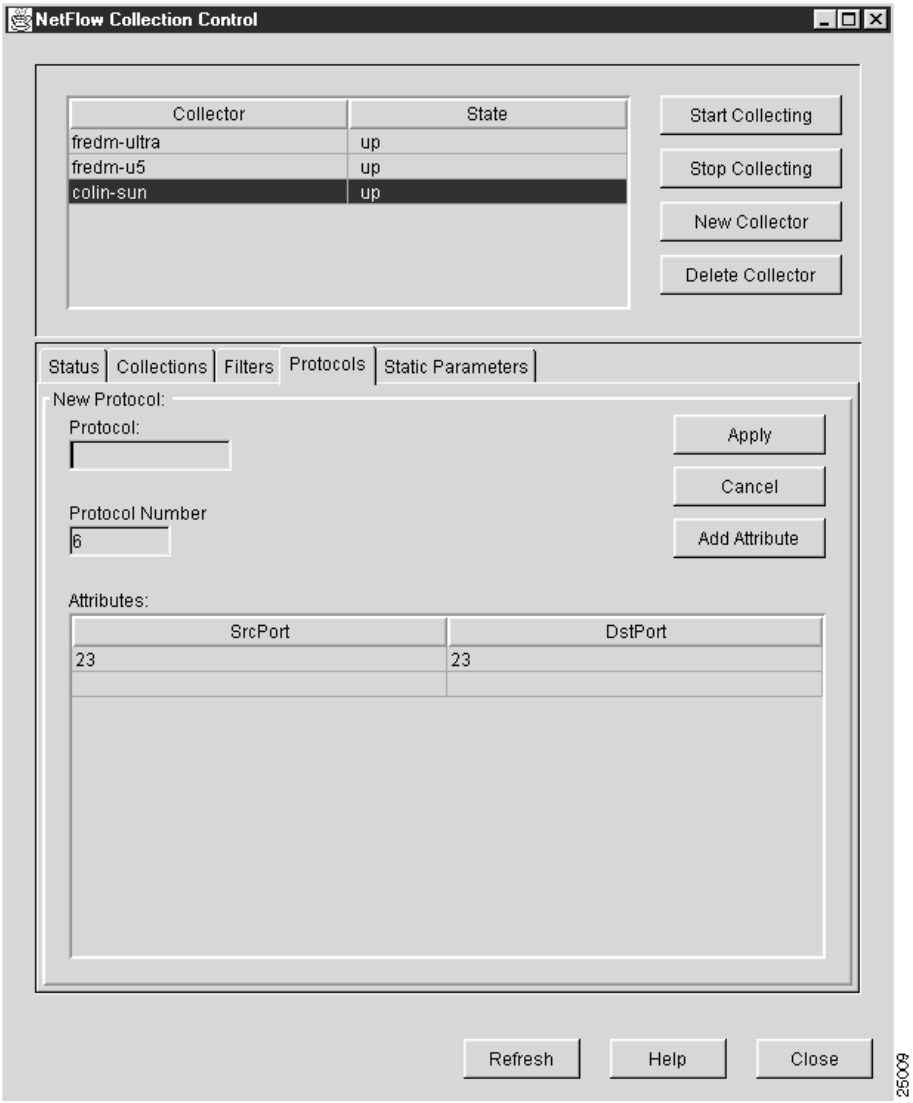

All of the parameters of the most recently selected Protocol panel are propagated into the New Protocols panel, giving you a basis for:

- Accepting the existing protocol parameters and associating them with a new protocol name.
- Changing the existing protocol parameters and associating them with a new protocol name.

In either case, the New Protocol panel incorporates all of the configuration facilities available to you by means of the Protocols panel. In addition, the New Protocol panel incorporates a Protocol pane through which you can define a unique name for the new protocol.

When you create a new protocol for the selected FlowCollector, a command is sent to the FlowCollector, causing it to update its nfknown.protocols file and to aggregate NetFlow data according to the new protocol.

The New Protocol panel provides the following new protocol definition facilities:

- Protocol ID text field—This text field enables you to create a unique alphanumeric name of up to 14 characters in length (with no intervening spaces) for a new protocol for the selected FlowCollector.
- Protocol Number text field—This field enables you to specify a protocol number (in the range from 1 to 255) for the new protocol named by means of the Protocol text field above.

  For example, TCP is always associated with protocol number 6; similarly, Telnet is always associated with protocol number 23.

  Any protocol number that you specify in this text field must agree with that assigned by the Internet Assigned Numbers Authority (IANA). The protocol number should be the same for all protocols of the same type (for example, TCP-TELNET and TCP-FTP are both protocol 6).
- Attributes pane—This pane enables you to specify the port numbers of the source and destination devices that are communicating by means of the selected protocol. Network devices communicate with each other by means of selected well-known protocols and selected well-known ports.

  The port numbers that you specify in the Attributes pane have a bearing on the way the FlowCollector translates and stores NetFlow data files. For example, if the new protocol is identified as protocol “6” (for TCP) in the Protocol Number field and the SrcPort and the DestPort are using port “20” (for FTP), the data files will carry the notation TCP-FTP when the FlowCollector writes the data files for the traffic flow. Thus, the port numbers that you specify in the Attributes pane are mapped to the new protocol name.

  You can employ multiple network protocols that have the same protocol number. These protocols, however, are associated with different ports, depending on the nature of the communicating applications.

  In general, port numbers are associated with application-level protocols. In all cases, exercise extreme care in modifying any existing pre-determined port number.
- Apply button—When you click Apply, the new protocol name is added to the list of protocols in the Protocol pane.

  Simultaneously, a command is sent to the selected FlowCollector, updating its nfknown.protocols file and causing it to recognize the new protocol definition.

Note If two or more entries in the nfknown.protocols file have the same protocol number and port number, the first entry encountered in this file is assumed to be valid and is used for mapping the translated name of the protocol into the FlowCollector’s stored NetFlow data files. Any duplicate entry in this file is ignored.

- Cancel button—Click Cancel to abandon an in-process protocol definition task.
- Add Attribute button—When you click Add Attribute, a row is added to the Attributes table, creating a placeholder for defining new source and destination port numbers.

Using the Static Parameters Panel

When you click the Static Parameters tab of the NetFlow Collection Control window, the Static Parameters panel appears (see Figure 3-64).

Figure 3-64 Sample Static Parameters Panel of NetFlow Collection Control Window

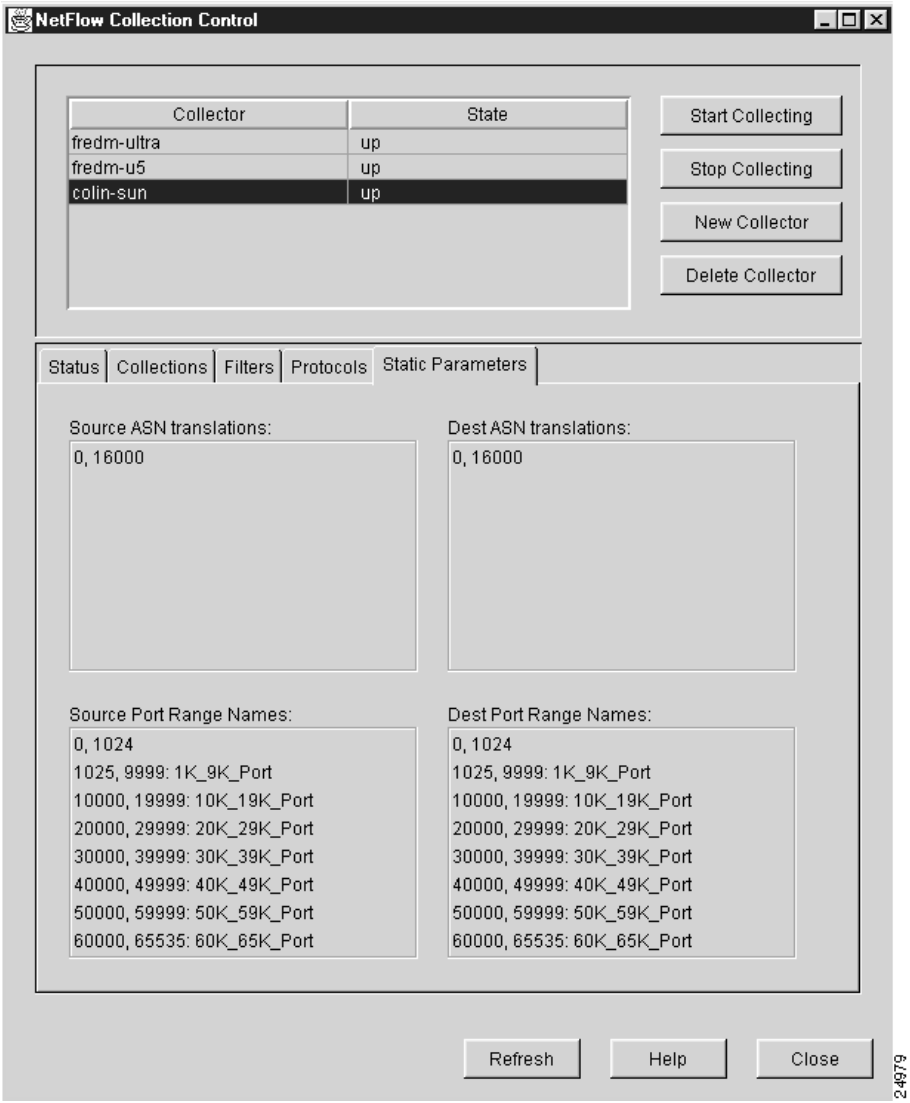

You cannot interact with the Static Parameters panel. It only displays static configuration information for the selected FlowCollector.

The Static Parameters panel has four panes, each of which shows the default text strings that are written into NetFlow data files in place of numeric information. These default settings are provided for your use in translating less readable numeric data into equivalent more readable text strings when NetFlow data files are stored by the selected FlowCollector.

A set of default static parameters and associated text strings are provided in the configuration files that are shipped with the FlowCollector, enabling you to take immediate advantage of this built-in translation capability.

Each of the four panes in the Static Parameters panel contains the following information:

- Numeric data—Each line of every pane begins with a numeric entry followed by a colon (:).
- Text string—Opposite every numeric entry is a text string that is written into the NetFlow data files for the selected FlowCollector.

The panes in the Static Parameters panel are described briefly below. If you wish to change any default setting, you must edit the appropriate FlowCollector configuration file, as identified below.

- Source ASN Translations pane—This scroll pane displays the source autonomous system numbers (ASNs) and corresponding text strings found in the following configuration file for the selected FlowCollector:

  /opc/CSCOnfc/config/nfknown.srcasns

- Dest ASN translations pane—This scroll pane displays the destination ASN numbers and corresponding text strings found in the following configuration file for the selected FlowCollector:

  /opt/CSCOnfc/config/nfknown.dstasns

- Source Port Range Names pane—This scroll pane displays the source port ranges and corresponding text strings found in the following configuration file for the selected FlowCollector:

  /opt/CSCOnfc/config/nfknown.srcports

- Dest Port Range Names pane—This scroll pane displays the destination port ranges and corresponding text strings found in the following configuration file for the selected FlowCollector:

  /opt/CSCOnfc/config/nfknown.dstports

Configuring Routers for Data Export

When you select the Router Configuration option of the Tools menu, the Configuration Files Directory window (see Figure 3-65) prompts you to specify the name of a directory on a network device that contains configuration files for the routers and switches in your network.

Figure 3-65 Configuration Files Directory Window

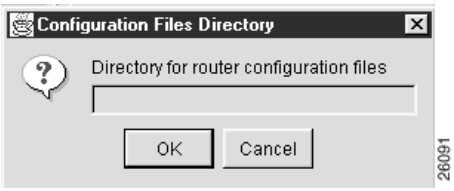

Creating Router Configuration Files

You create router configuration files at installation time when you configure devices to operate in the network. Typically, you log on to a newly-installed device and copy its running configuration file into a directory that you create on a UtilityServer host for storing router configuration files.

There must be a configuration file in this directory for each NetFlow or TMS export-capable device that you intend to configure by means of the Router Configuration window (see Figure 3-66).

Figure 3-66 Router Configuration Window

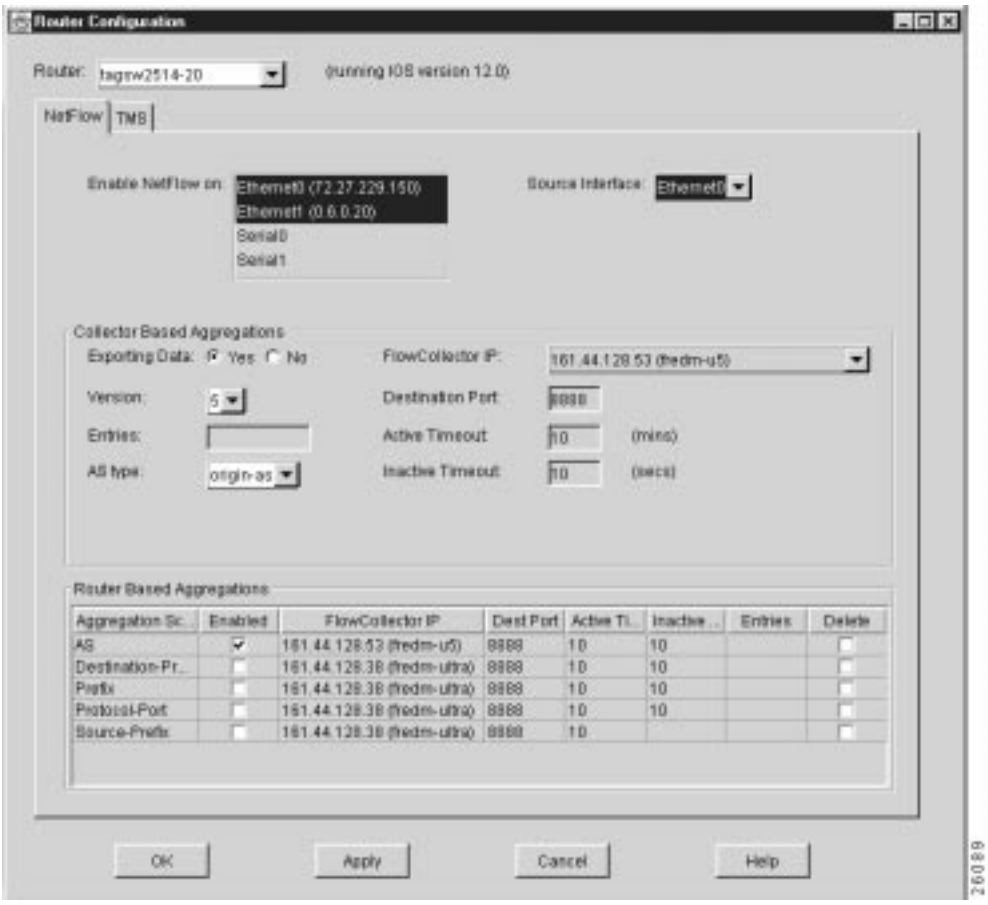

The configuration file directory provides the names, interfaces, and passwords of network devices, enabling you to log on to any NetFlow or TMS export-capable device in the network and configure it by means of the Router Configuration window.

Although you typically store router configuration files on a UtilityServer host, you can store them on any NFS-mounted volume in the network that is accessible to the UtilityServer. Thus, the directory path name that you enter in the Configuration Files Directory window can point to any UNIX directory on any host in the network that is reachable by the UtilityServer.

For example, assume that you have created a configuration file directory named "/tftpboot/configs/" on a UtilityServer operating on a UNIX platform. After you enter this directory path name (or any other directory path name that points to router configuration files), click OK.

If the directory path name that you enter is valid, the Router Configuration window appears (see Figure 3-66) with the NetFlow panel selected by default. The Router pull-down selection box incorporated into this window enables you to select any NetFlow or TMS export-capable device of interest and to configure that device for data export.

Elements Common to NetFlow and TMS Configuration Panels

The following facilities are common to the NetFlow configuration panel and the TMS configuration panel of the Router Configuration window:

- Router pull-down selection box—This single selection facility in the Router Configuration window enables you to display a list all of the network devices for which a valid configuration file exists. By default, the first router encountered in the configuration files directory appears as the selected entry in the Router selection box.

  When you select a router from the pull-down list, the Router Configuration window is refreshed to display the configuration information currently in effect for the selected device.

- OK push-button—When you click OK, the configuration parameters in the Router Configuration window are applied to the currently selected device, returning you to the main Display module window.

- Apply push-button—When you click Apply, any changes that you have made to the configuration parameters in the Router Configuration window are applied to the currently selected device.

- Cancel push-button—When you click Cancel, the Router Configuration window is closed without effecting any of the configuration parameters contained therein. You are returned you to the main Display module window.

- Help push-button—When you click Help, the Analyzer's context-sensitive online help system is made available to you.

Facilities of the NetFlow Configuration Panel

The Router Configuration window appears with the NetFlow tab selected by default, displaying the NetFlow configuration panel shown in Figure 3-66. This panel enables you to configure network devices to export NetFlow data.

The NetFlow configuration panel is partitioned into three functional areas, as described below:

- Interfaces area—This area at the top of the NetFlow configuration panel is used to configure router interfaces to support the generation and export of NetFlow records.

- Collector Based Aggregations—This area of the NetFlow configuration panel is used exclusively to configure routers to support FlowCollector-based aggregation of data from Version 1 and Version 5 export records.

- Router Based Aggregations—This area of the NetFlow configuration panel is used exclusively to configure routers to support router-based aggregation of Version 8 export records.

  Note that the first column of the table for the router based aggregations identifies the five on-router aggregation schemes used by the selected router to summarize Version 8 export records:

  — AS

  — Destination-Prefix

  — Prefix

  — Protocol-Port

  — Source-Prefix

The Collector Based Aggregations area and the Router Based Aggregations area of the Router Configuration window are functionally equivalent, the only difference being that each area enables you to control a specific type of NetFlow data collection and aggregation process.

You can use either area of the Router Configuration window at any time to configure selected routers to export or process NetFlow data as desired, including configuring the same router to perform the following tasks simultaneously:

- Export Version 1 or Version 5 records to a specified FlowCollector for processing and storage of NetFlow data according to a selected aggregation scheme (FlowCollector-based data aggregation).
- Process Version 8 export records on the selected router according to a selected on-router aggregation scheme for delivery to a specified FlowCollector (router-based data aggregation).

Configuration Parameters for Interfaces

The selection lists for configuring the interfaces on the selected router are described below:

- Enable NetFlow on—This multiple-selection list allows you to enable or disable NetFlow switching on each of the listed router interfaces. If you select an interface, NetFlow switching is enabled on that interface.

  NetFlow switching should be enabled on at least one interface in order for a FlowCollector-based or a router-based aggregation scheme to take effect.

- Source Interface—This single-selection pull-down list enables you to choose the source interface IP address to be used in the NetFlow export datagram. The IP address of this interface is used as the packet's source address. If you select “None,” the router will select the IP address of one of its interfaces.

The configuration parameters for the router interfaces apply to both FlowCollector-based data aggregations and router-based data aggregations.

Configuration Parameters for FlowCollector-Based Aggregations

The data fields and push-buttons for controlling FlowCollector-based aggregations (Version 1 or Version 5 export records) for a selected router are described below:

- Exporting Data radio buttons—These buttons enable you to control the export of data from the selected device to a specified FlowCollector.

  When you click Yes, the selected router begins sending NetFlow data to the specified FlowCollector (see the description of the FlowCollector IP field below).

  When you click No, the selected router stops sending NetFlow data to the specified FlowCollector.

- Version pull-down box—This single selection box enables you to select the export record type (either Version 1 or Version 5) that you want the selected router to use in exporting data to the specified Flowcollector.

  Click on the down arrow in this field and drag the mouse pointer to select the desired export record type.

- Interface pull-down box—This single selection box enables you to select the interface that the selected router is to use in exporting NetFlow data to the specified FlowCollector.
- Entries field—This field enables you to specify the maximum number of entries to be accumulated in the main cache of the selected router.
- AS type pull-down box—This single selection box enables you to specify the AS (autonomous system) type that the selected router looks for in monitoring the flow of traffic through the selected router (valid for Version 5 export record type only).

The possible AS types are:

— Origin-AS—Causes the selected router to monitor an end-to-end traffic flow.

In the origin AS case, the selected router concerns itself with the traffic passing between the source and the destination devices in the overall traffic flow.

— Peer AS—Causes the selected router to monitor an intermediate traffic flow.

In the peer AS case, the selected router concerns itself only with the data passing between itself and the routers on either side.

- FlowCollector IP pull-down box—This single selection box enables you to display a list of the FlowCollectors currently configured to operate in your network and to select one of them as the target host to which the NetFlow export records from the selected router are to be sent.

  To select a FlowCollector from the pull-down list, hold down the mouse button on the down arrow in the box, drag the mouse pointer to the desired FlowCollector, and release the mouse button.

  Alternatively, if the desired FlowCollector is not defined, an IP address can be entered into the text field.

  The IP address and the name of the selected FlowCollector are then propagated into the FlowCollector IP selection box, establishing the selected FlowCollector as the target host for storing NetFlow data.

  Note that if you select the “No” option for the Export Data radio button (described above) and click Apply (thus altering the export state of the selected router), the following message is displayed.

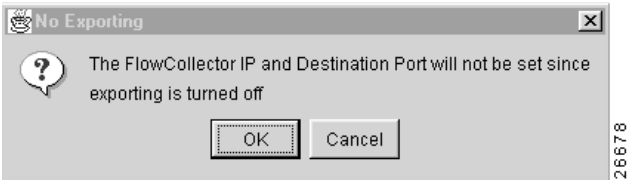

  When you click OK in this message box, the window closes, and the FlowCollector IP box, as well as the Destination Port box (see next bullet), are cleared.

  When you click Cancel in this message box, you are returned to the NetFlow configuration panel without further action being taken.

- Destination Port box—Enables you to specify the port on the selected FlowCollector to which the exported NetFlow data is to be sent.
- Active Timeout field—Enables you to specify the maximum number of minutes that a flow can remain in the aggregation cache of the router. The default value is 15 minutes.
- Inactive Timeout field—Enables you to specify the number of seconds that an inactive flow can remain in aggregation cache before being aged out. The default value is 30 seconds.

Note You must click Apply in the Router Configuration window to put into effect any changes that you make to any of the FlowCollector based configuration parameters described above.

Configuration Parameters for Router-Based Aggregations

In general, the column headings in the table for router-based aggregations (Version 8 export records) parallel the functions of the five similarly-named data fields described in the preceding section for FlowCollector-based aggregations.

Beyond the first column (which lists the five router based aggregation schemes), all columns of the table are editable. For example, you can click any cell in the table at any time, change the data contained therein to a desired value, and click Apply to effect the change.

The functions of the columns in the table are described below in their order of appearance from left to right:

- Enabled column check boxes—Enables you to select the desired aggregation scheme(s) by which traffic data is to be summarized before being sent to the specified FlowCollector (see the description of the FlowCollector IP column below).
- FlowCollector IP column—Enables you to select from a list of known FlowCollectors the one to which you want the selected aggregation scheme data to be sent.

  When you click on any row (cell) in this column, a pull-down list of currently known FlowCollectors is displayed, enabling you to drag the mouse to any desired FlowCollector and select it as the target for receiving the aggregation scheme data. Or, you can enter a desired IP address if it is not in the current list.
- Dest Port column —Enables you to specify the port on the selected FlowCollector to which the selected aggregation scheme data is to be sent.
- Active Timeout column—Enables you to specify the maximum number of minutes that a flow can remain in the aggregation cache. The default value is 15 minutes.
- Inactive Timeout column—Enables you to specify the maximum number of seconds that an inactive flow can remain in aggregation cache before being aged out. The default value is 30 seconds.
- Entries column—Enables you to specify the maximum number of entries that can be accumulated in the aggregation cache of the selected router.
- Delete column check boxes—Enables you to delete any one, some, or all of the currently selected aggregation schemes for the selected router.

Note You must click Apply in the Router Configuration window to put into effect any changes that you make to any of the router based configuration parameters described above.

Facilities of the TMS Router Configuration Panel

When you select the TMS tab in the Router Configuration window, the TMS router configuration panel shown in Figure 3-67 is displayed. This panel enables you to configure network devices to export TMS data.

Figure 3-67 Window for Configuring TMS Devices

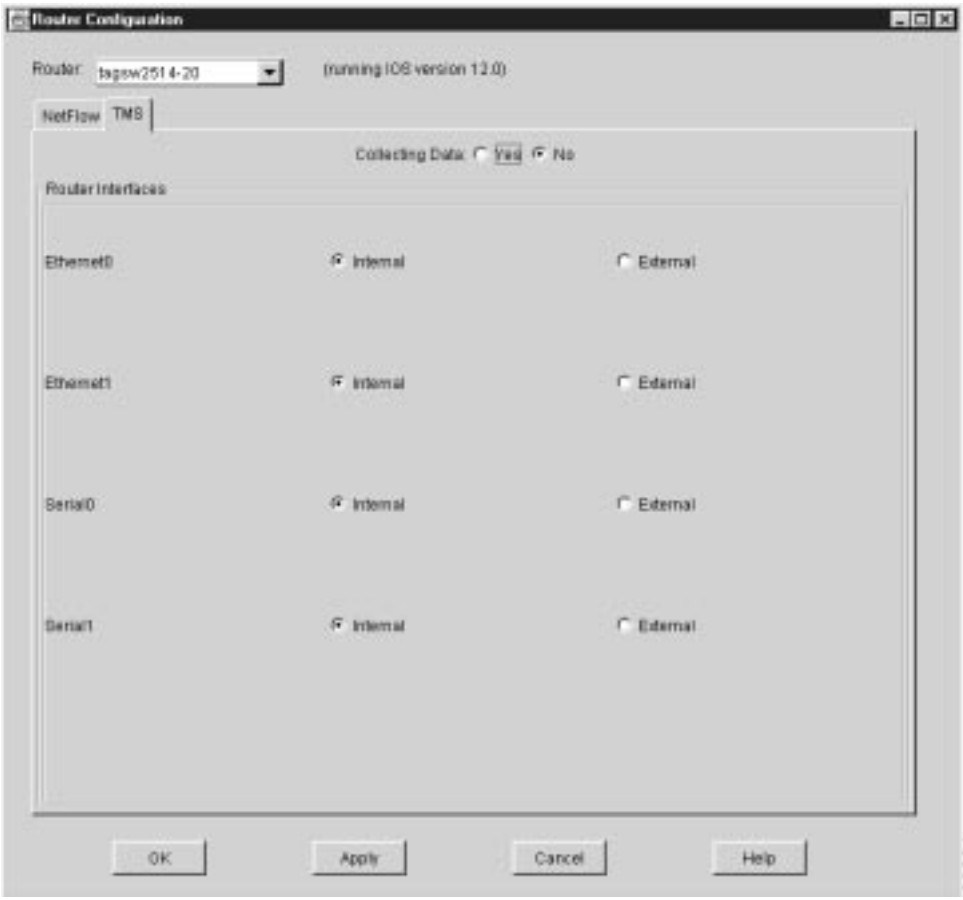

The following facilities displayed in conjunction with the TMS router configuration panel perform the same functions as those described in the "Elements Common to NetFlow and TMS Configuration Panels" section on page 3-105:

- Router pull-down selection box
- OK push-button
- Apply push-button
- Cancel push-button
- Help push-button

The functional elements unique to the TMS router configuration panel include the following:

- Router interfaces list—For each router interface listed in the TMS router configuration panel, you can specify whether it is an internal or an external interface.

  Specifying internal or external determines the form in which the TMS data is displayed, as described in the "TMS TrafficMatrix Data Aggregation Scheme" section on page 3-8 of this chapter.

- Collecting Data radio buttons—These single selection buttons enable you to start or stop incrementing counters in the selected router.

When you click Yes, TMS collections on the selected router are started; when you click No, TMS collections on the selected router are stopped.

- Internal/External radio buttons—These single selection buttons enable you to designate the interface as being associated with either “internal” data or “external” data. Only one radio button for an interface can be selected at one time.

  You can designate any combination of internal or external states among the listed interfaces, that is, some interfaces can be internal, while others can be external, but you must select either one state or the other for each listed interface.

Router Configuration Error Conditions and Messages

In the event that the directory path name that you specify in the Configuration Files Directory window (see Figure 3-65) is invalid or does not exist, the following Router Config Files window is displayed by the Analyzer to so indicate.

In this case, click OK to return to the main Display module window.

You can then start over by selecting the Router Configuration option of the Tools menu. When the Configuration Files Directory dialog window again appears, you can correctly specify the path name of the desired configuration files directory, or you can specify the name of a different directory that you have reason to believe contains requisite router configuration files.

Alternatively, you can open a UNIX window and, by means of CLI commands, copy the router configuration files from their known current directory into the directory that you defined by means of the Configuration Files Directory dialog window.

In any case, you must have access to a directory that contains valid router configuration files before you can open and use the Router Configuration window. This window can only display information pertaining to network devices for which valid configuration files exist.

When the Router Configuration window appears, you can select any NetFlow or TMS data exporting device listed in the window’s Router pull-down list and view and/or change any of its associated configuration parameters.

The UtilityServer checks all the configuration files stored in the router configuration file directory, ignoring any invalid files. After the UtilityServer validates the router name, the user’s login password, and the enable password for each router, it passes the router name and its associated configuration parameters to the Display module. The name of the router is then added to the Router pull-down list, and, if selected, the data fields of the Router Configuration window are updated accordingly.

If for some reason a connection cannot be made to a device that you select in the Router pull-down list, the following No Connection message window appears to so indicate.

In this case, click OK to close the window. You can then attempt to connect to any other router listed in the Router pull-down list of the Router Configuration window.

Hints Menu Option

The Display module Hints menu provides a Get Hints option. When you select this option, the pop-up window shown in Figure 3-68 appears. You can access this window at any time for assistance in performing the common Analyzer tasks listed in the window.

Figure 3-68 Typical Hints Pop-up Window

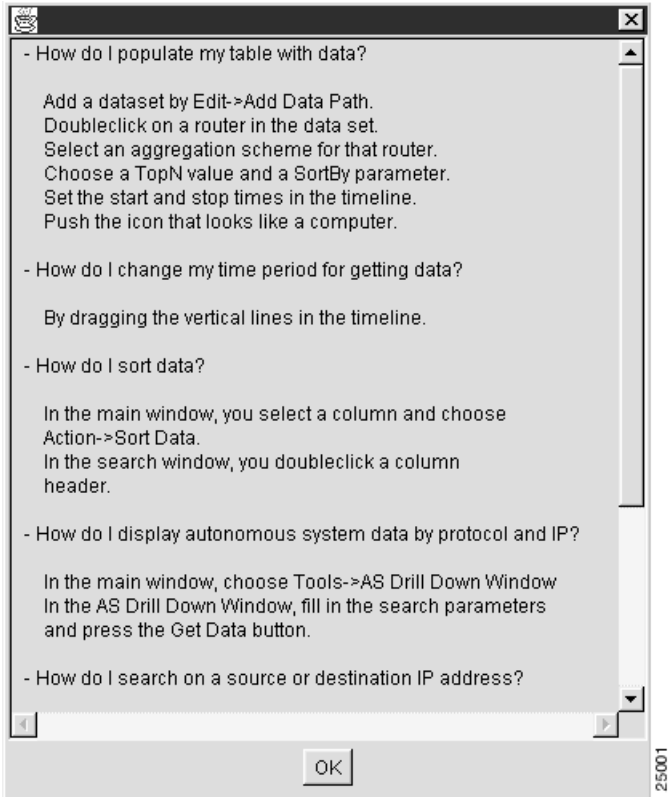

Help Menu Options

The Analyzer Display module Help menu provides the selectable options described in the following sections.

About Network Data Analyzer

Select the About Network Data Analyzer option of the Tools menu to invoke Netscape Navigator. The resulting window, see Figure 3-69, lists the platforms on which the Analyzer is supported and tells you how to obtain assistance.

Figure 3-69 Netscape "About NetFlow FlowAnalyzer" Window

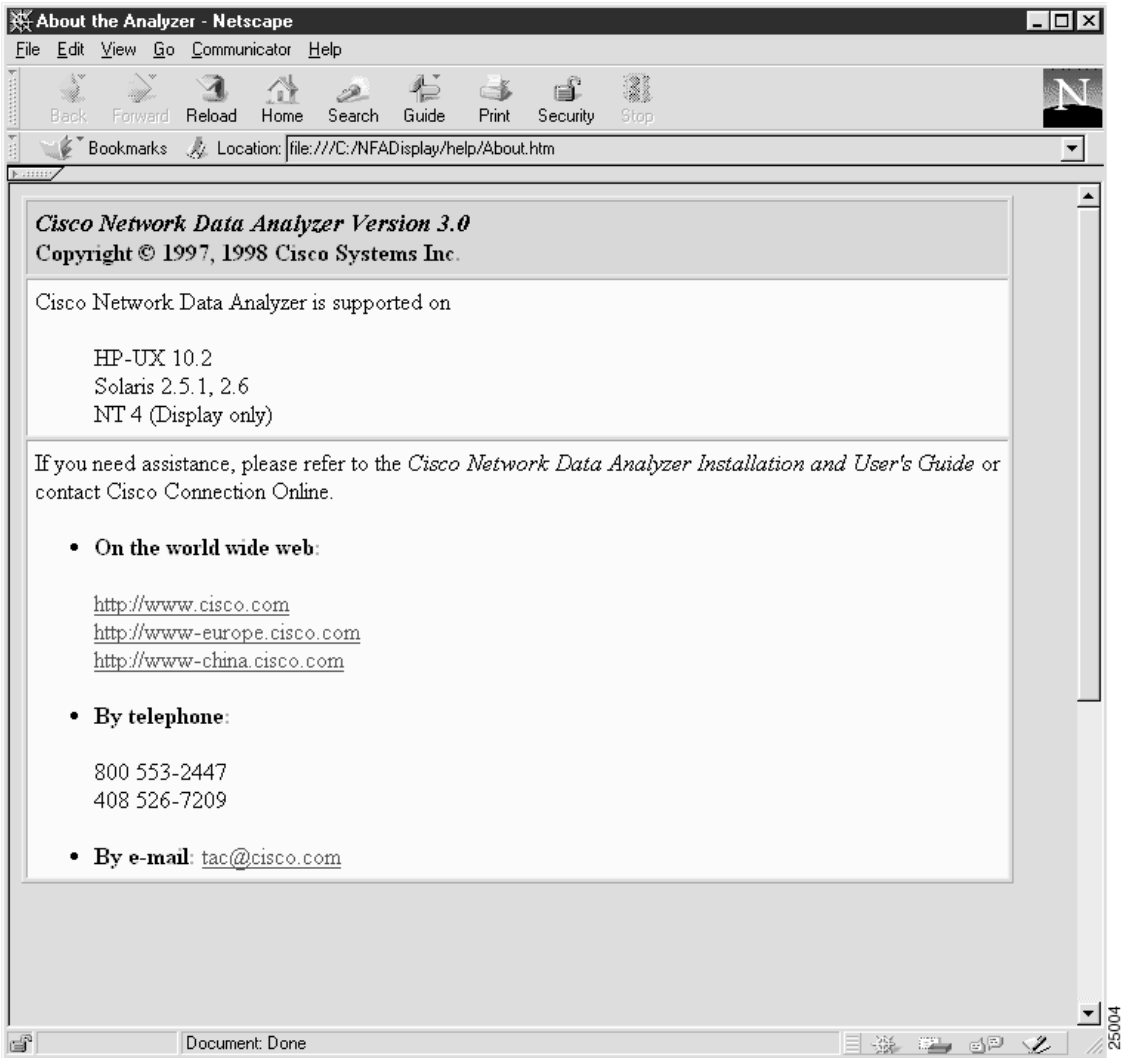

How to Use the Network Data Analyzer

Select the How to Use the Network Data Analyzer option of the Help menu to invoke Netscape Navigator. The resulting window, see Figure 3-70, lists typical and optional Analyzer tasks that you can select at random for information about how to perform the selected task.

Figure 3-70 Netscape "How to Use the Network Data Analyzer" Window

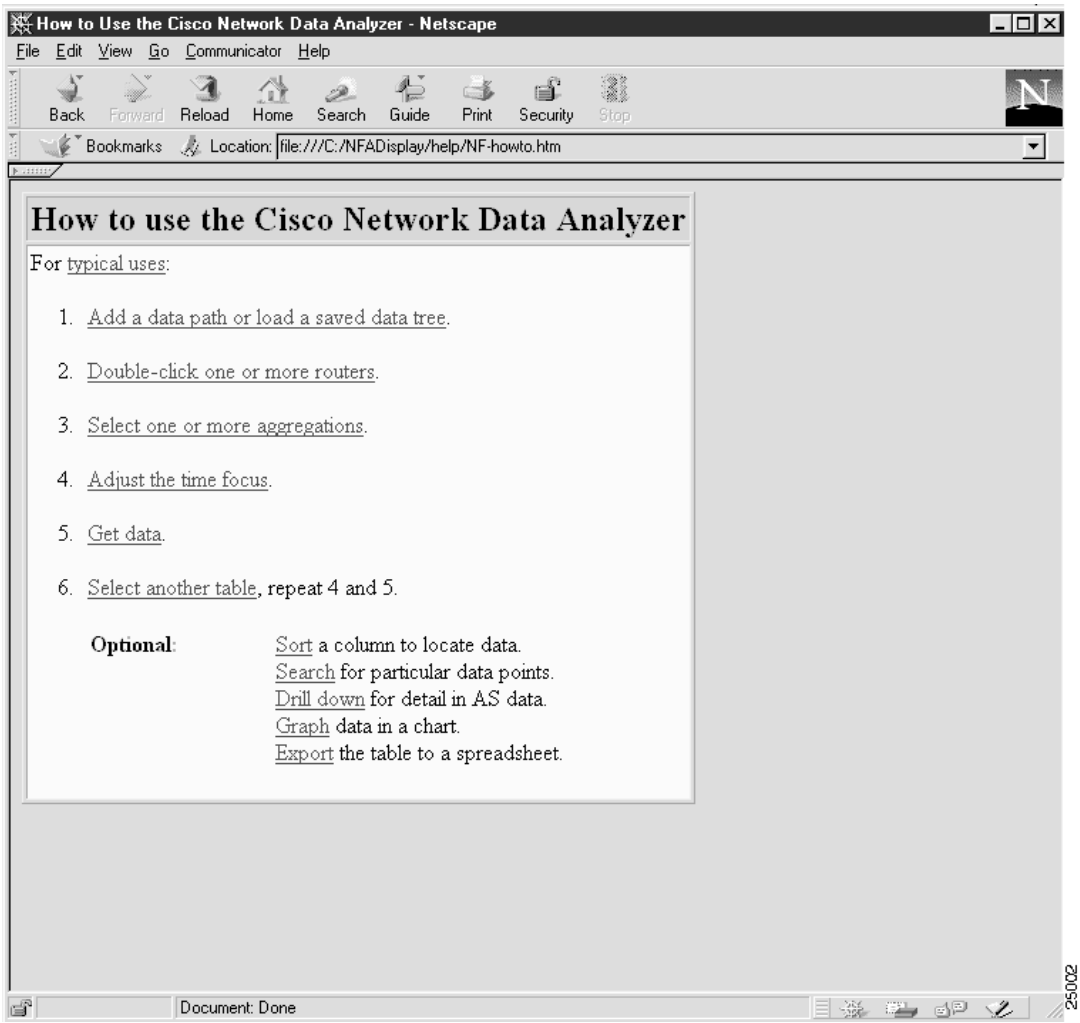

Tables and Aggregation Schemes

Select the Tables and Aggregation Schemes option of the Help menu to invoke Netscape Navigator, which displays the window shown in Figure 3-71. The table in this window lists the aggregation schemes available for use with the Analyzer.

You can select any entry in this table at random for additional information about the characteristics, uses, and parameters of the selected aggregation scheme.

Figure 3-71 NetScape “Data Tables” Window

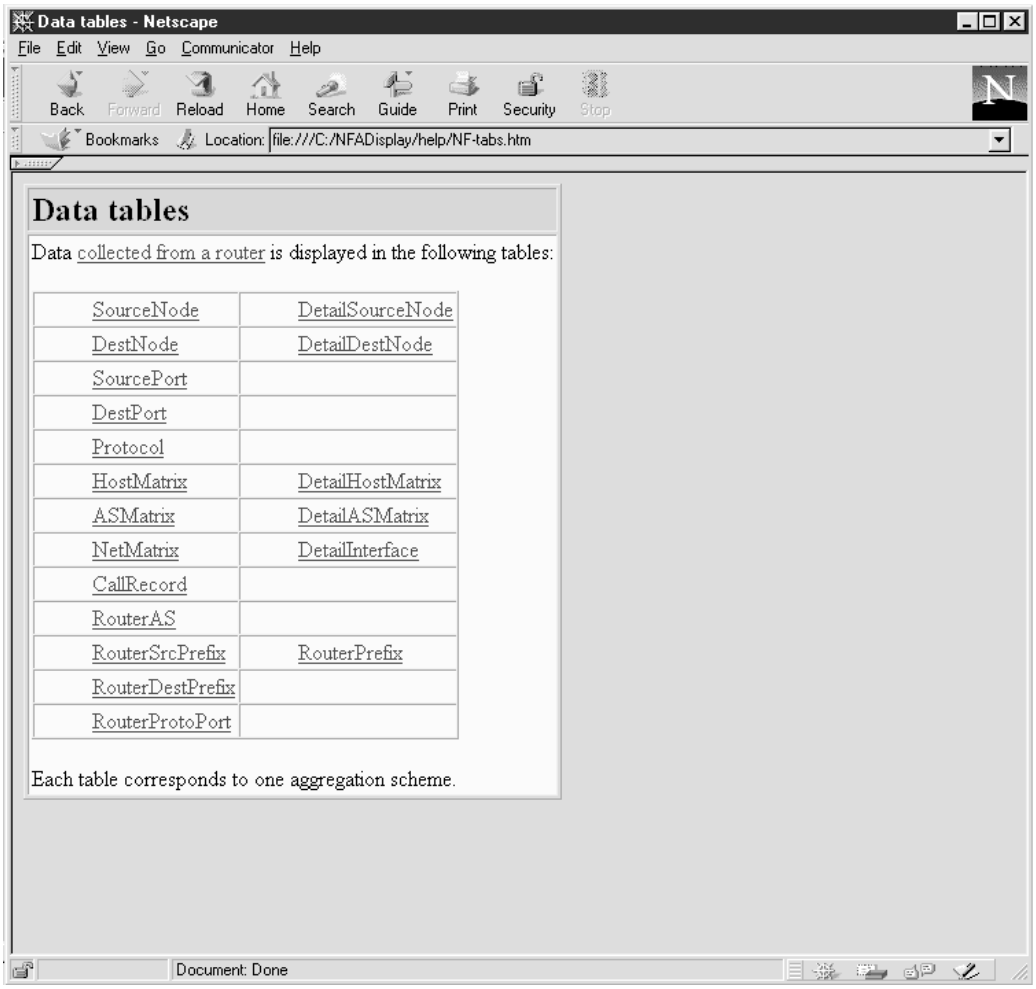

Getting Support

Select the Getting Support option of the Help menu to invoke Netscape Navigator, which displays the window shown in Figure 3-72. This window lists the platforms on which the Analyzer runs and the resources that you can call upon for assistance in using the Analyzer.

Figure 3-72 NetScape “Getting Support” Window