Overview of the NetFlow FlowAnalyzer

NetFlow FlowAnalyzer Version 2.0 is a network analysis tool that you can use to display and analyze network traffic information collected from Cisco NetFlow-enabled devices running NetFlow services software.

NetFlow FlowAnalyzer Version 2.0 requires the use of Cisco IOS Release 11.1CC or 11.1(11)CA (or later) running on Cisco NetFlow-enabled devices.

Cisco IOS software provides enhanced security, quality of service, and traffic management capabilities in a NetFlow services environment. Related network services are based on traffic flows between network layer address pairs that are visible to the NetFlow FlowCollector application.

Important Information Regarding This Product

NetFlow FlowAnalyzer Version 2.0 has been expressly designed to operate with NetFlow FlowCollector Version 2.0. However, for the benefit of users of previous versions of these applications, it is important to note the following:

- Supported configurations:
  - FlowAnalyzer Version 2.0 processing data collected by FlowCollector Version 2.0.
  - FlowAnalyzer Version 2.0 processing data collected by FlowCollector Version 1.0.
- Unsupported configuration:
  - FlowAnalyzer Version 1.0 processing data collected by FlowCollector Version 2.0; FlowAnalyzer Version 1.0 can only process data collected by FlowCollector Version 1.0.
- Constraints/requirements:
  - If you use FlowAnalyzer Version 2.0 to process data collected by FlowCollector Version 1.0, you cannot get data for the DetailASMatrix aggregation scheme.

    The DetailASMatrix aggregation scheme is new to both FlowAnalyzer Version 2.0 and FlowCollector Version 2.0.
  - FlowAnalyzer Version 2.0 requires the use of the DetailASMatrix aggregation scheme for display functions in the following windows:

    (a) The AS Drill Down window (see the section entitled “Drilling Down on Network Flows” in Chapter 3).

    (b) The Search Window (see the section entitled “Searching for Flows by Source and Destination Addresses” in Chapter 3).

NetFlow Services

NetFlow services provide high-performance, network-layer switching that enables you to capture, display, and analyze a rich set of statistics gathered from network traffic flows. Such information includes:

- Per-user statistics
- Per-protocol statistics
- Per-port statistics
- Per-type-of-service statistics

NetFlow data being exported by NetFlow-enabled devices in your network to host FlowCollector workstations can be retrieved and used for a variety of purposes, such as the following:

- Network planning, management, and analysis
- Network load balancing
- Internet service provider (ISP) billing
- Data warehousing/mining for marketing purposes

NetFlow services also facilitate the processing of security access lists. It is much more efficient in this regard than other switching methods.

In both NetFlow services and conventional switching services, the first packet in a traffic flow undergoes fairly complex processing that involves switching, security, queuing, and accounting tasks.

In conventional switching, every subsequent packet in a traffic flow undergoes the same complex processing as the first packet. However, in NetFlow services, subsequent packets in a traffic flow are processed faster due to the inherent efficiency of the NetFlow services software.

NetFlow System Elements

The NetFlow system (see Figure 1-1) enables you to retrieve and analyze NetFlow data that has been collected from NetFlow-enabled devices in your network. The FlowCollector application, running on a designated workstation in your network, is the means by which the NetFlow information is collected and made available for use by the FlowAnalyzer Display module.

The FlowCollector stores this information in a locally managed database that can be accessed on-demand by means of user commands issued at the Display module console.

The NetFlow system has three main components:

- DataExporter function of Cisco IOS software—Native Cisco IOS software functionality, operating in conjunction with NetFlow-enabled devices in your network, sends NetFlow data to designated host FlowCollector workstations in your network. This NetFlow data is transmitted to FlowCollector workstations by means of UDP datagrams (see Figure 1-1).
- FlowCollector—The NetFlow FlowCollector, running on one or more designated workstations in your network, receives NetFlow data being exported by the DataExporter and stores the data locally as flat files in a hierarchically structured UNIX directory.
- FlowAnalyzer—The DisplayServer module of the NetFlow FlowAnalyzer application retrieves NetFlow data from any FlowCollector workstation in the network and transmits the data to the Display module of the NetFlow FlowAnalyzer application for presentation on the screen of a host workstation or PC.

The Display module of the FlowAnalyzer (see Figure 1-1) provides a graphical user interface (GUI) for the NetFlow system. This module enables you to select a data set (a collection of NetFlow traffic information) pertaining to a single router in the network or to multiple routers in the network and to display the information in a particular format (a defined aggregation scheme).

The FlowCollector application supports a variety of NetFlow data aggregation schemes, any one of which can be selected at any time to display applicable Netflow data for any router of interest.

You use the Netflow Data area of the Display module window (see Figure 3-4) to select any aggregation scheme currently applicable to a network router of interest. You can then issue appropriate Display module commands to initiate the display of relevant NetFlow data.

The NetFlow system thus provides a powerful tool for analyzing network traffic flows involving network routers. You can use the data gathered from such routers to better balance network loading, resolve network problems, and optimize network performance.

Figure 1-1 Elements of the NetFlow System

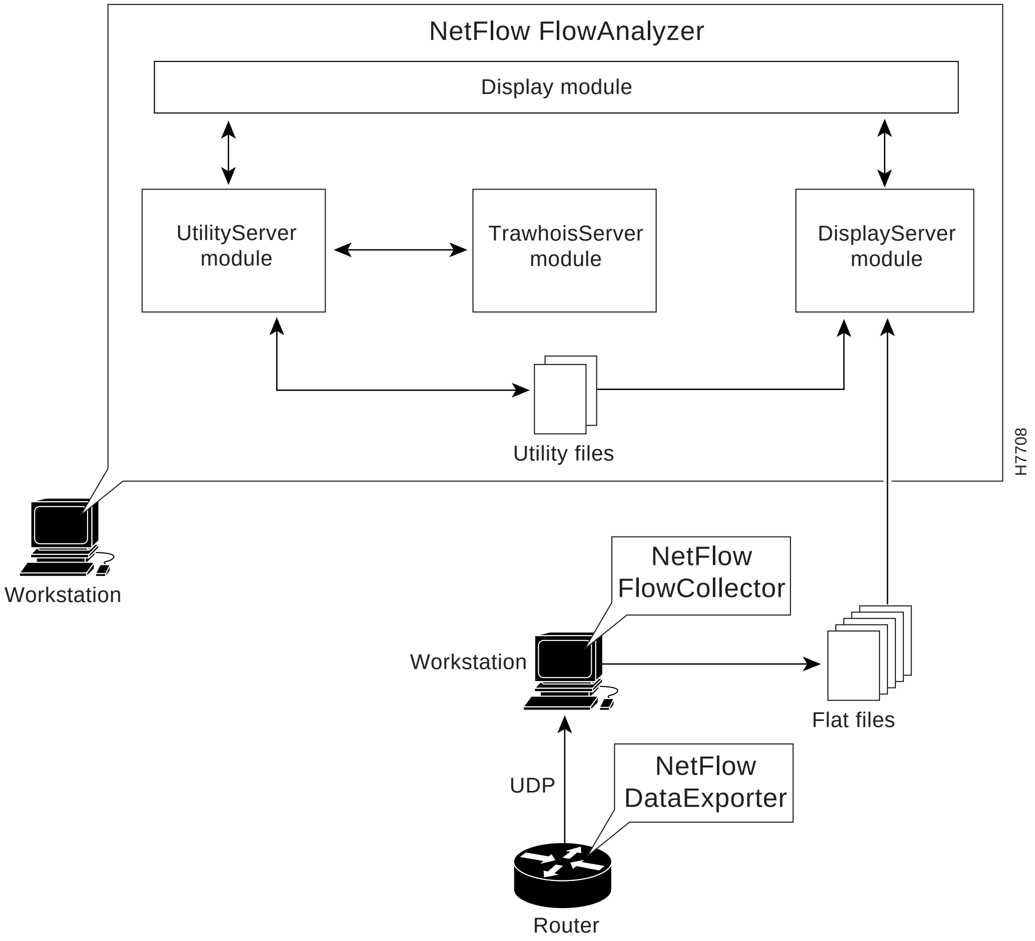

The components of the NetFlow system are described in greater detail in the following sections.

NetFlow DataExporter

The export of NetFlow data is accomplished by native functionality within Cisco IOS software that is running on NetFlow-enabled devices in your network. This functionality supports the bulk export of NetFlow traffic data to designated FlowCollector workstations in your network.

A summary of traffic statistics for all expired router traffic flows can be exported on a periodic basis (by means of UDP datagrams) to a designated workstation in your network that is running the FlowCollector application (see Figure 1-1).

The FlowCollector collects traffic data being exported from network routers and stores the data in a designated local directory. Hence, the primary role of the FlowCollector is to maintain a database of NetFlow data files. This database can be accessed on demand by the DisplayServer module to retrieve NetFlow data on behalf of the Display module. You then issue appropriate commands at your Display module console to display the requested NetFlow data.

NetFlow FlowCollector

The FlowCollector application supports fast, scalable, reliable, and economical collection of NetFlow data that is being exported from multiple NetFlow-enabled devices in your network.

The collected NetFlow data is stored on a designated FlowCollector workstation in the network for later retrieval, display, and analysis by means of the Display module of the FlowAnalyzer application.

Important FlowCollector functions include the following:

- Collecting NetFlow data from multiple NetFlow-enabled devices in your network
- Reducing the volume of collected NetFlow data through the use of filters and aggregation schemes
- Storing the collected NetFlow data in a user-defined directory on the host FlowCollector workstation
- Managing the local storage space on the FlowCollector workstation for the NetFlow data files

The following sections describe functions of the FlowCollector application in greater detail.

Traffic Statistics

NetFlow data records contain detailed traffic information for expired network traffic flows. Such information includes statistics about Layer 3 source and destination nodes in the network, down to the level of the individual applications and protocols participating in end-to-end communications within the network.

The ability to collect, display, and analyze detailed network traffic statistics provides the following benefits:

- Enables network managers to monitor network traffic and fine-tune network performance. By gathering and analyzing network traffic data, network managers can make better decisions with respect to such matters as the following:
    - — Which users and applications may need additional bandwidth
    - — Which users may require a specific quality of service
- Enables traffic information to be consolidated and used for billing purposes on a per-application or actual-usage basis.

Traffic statistics can include any one or all of the following types of information:

- Time stamp of data flow
- Source and destination IP addresses of the data flow
- Source and destination port numbers of communicating nodes
- Port numbers of network devices
- Next hop addresses
- Total bytes in a data flow
- Number of packets in a data flow
- First and last time stamps of packets switched as part of a data flow
- Sequence numbers
- Source and destination autonomous system (AS) numbers
- Source and destination prefix masks

NetFlow Data Aggregation Schemes

The FlowCollector application collects detailed traffic statistics and aggregates (summarizes) the data according to one or more of the following aggregation schemes:

- SourceNode
- DestNode
- HostMatrix
- Protocol
- SourcePort
- DestPort
- DetailDestNode
- DetailHostMatrix
- DetailInterface
- CallRecord
- ASMatrix
- DetailASMatrix
- DetailSourceNode
- NetMatrix

For detailed descriptions of the NetFlow data aggregation schemes, see the section entitled “Drilling Down on Network Flows” in Chapter 3.

Filters

To customize traffic statistics for display and analysis purposes, you can apply the following types of filters to the NetFlow data being stored by the FlowCollector application:

- Srcaddr —Network layer IP address of the source node
- Dstaddr —Network layer IP address of the destination node
- Nexthop router IP address
- Srcinterface—Interface number of input device (physical interface)
- Dstinterface—Interface number of output device (physical interface)
- Srcport number—Transport layer port number of the source node (per RFC 1700)
- Dstport number—Transport layer port number of the destination node (per RFC 1700)
- Src AS Number
- Dst AS Number
- Type of Service (ToS)—ToS byte from IP header
- NetFlow Export Datagram Source IP address

Typical FlowCollector Topology

Figure 1-2 shows a typical topology for a network that incorporates NetFlow-enabled devices that are exporting Netflow data to a designated FlowCollector workstation in the network.

NetFlow-enabled devices deliver NetFlow data to the IP address and UDP port number of designated FlowCollector workstations in your network.

After you configure and start the FlowCollector application, it listens to the UDP ports of the NetFlow-enabled devices in your network to collect the traffic data being exported by such devices by means of UDP datagrams.

Figure 1-2 Typical FlowCollector Topology

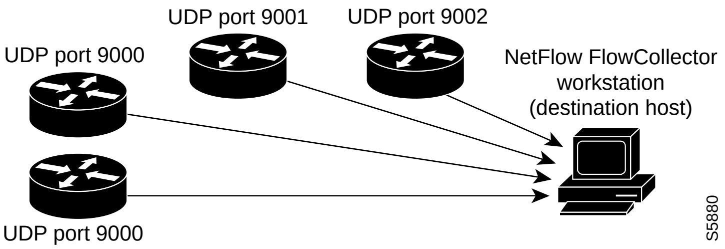

NetFlow FlowAnalyzer

The FlowAnalyzer application is a network analysis tool that combines a graphical user interface with other companion modules. Together, these components enable you to retrieve, display, and analyze NetFlow data that has been collected from specified NetFlow-enabled devices in your network.

The individual modules of the FlowAnalyzer application are described in the following sections.

Display Module

The Display module is a stand-alone Java application that can be installed and run on workstations or PCs in your network. This module provides an easy-to-use, graphically oriented user interface to the NetFlow system. You can install any number of Display modules into your network, as appropriate for your particular NetFlow data retrieval and analysis requirements.

The Display module can present NetFlow data in a variety of formats, including the following:

- Tables
- Graphs—Pie charts, bar charts, or histogram charts
- Numeric lists
- Non-numeric lists

In addition, the Display module incorporates a number of facilities that help you to use the FlowAnalyzer application conveniently and effectively. Such user aids include the following:

- Pop-up windows—Several pop-up windows have been incorporated into the Display module user interface to provide helpful tips and reminders in using the FlowAnalyzer.

  One such window is shown in Figure 1-3. You have the option to suppress a pop-up window whenever it appears if you are familiar with the task to which it applies.

- Movable time slider marks (see Figure 1-3)—These two vertical bars in the upper right area of the Display module window enable you to establish the applicable time interval for FlowAnalyzer data retrieval and analysis tasks.

Figure 1-3 User Aids Incorporated into the Display Module Interface

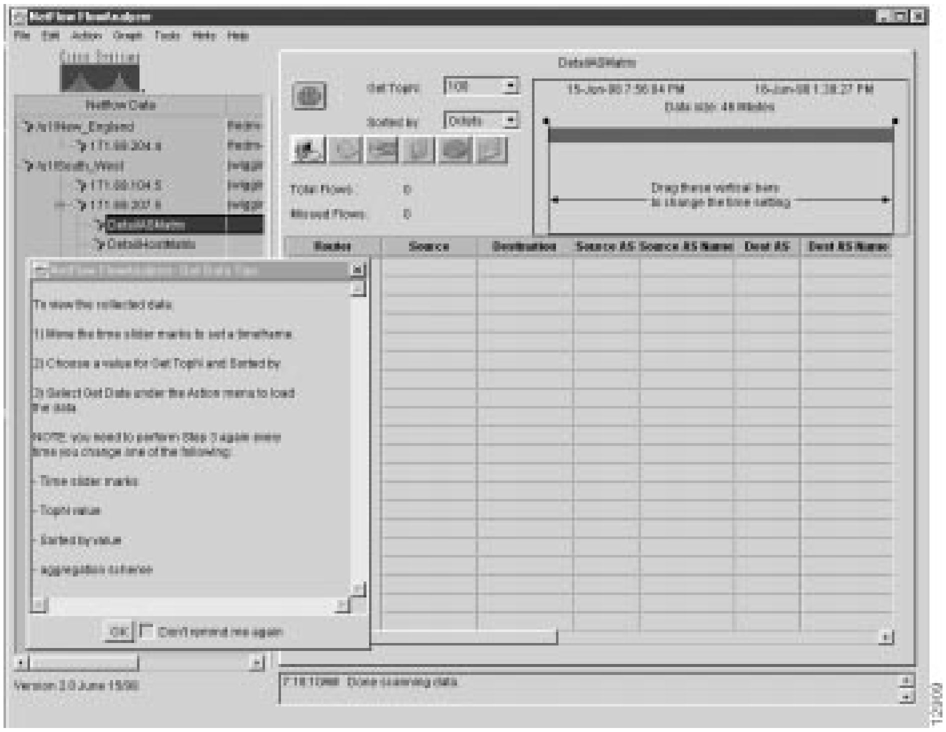

To view the collected data:

1) Move the time slider marks to set a timeframe

2) Choose a value for Get TopN and Sorted by

3) Select Get Data under the Action menu to load the data

NOTE: you need to perform Step 3 again every time you change one of the following:

- Time slider marks
- TopN value
- Sorted by value
- aggregation scheme

Get TopN: 100

Sorted by: Octets

15-Jun-98 7:56:34 PM

16-Jun-98 1:38:27 PM

Data size: 48 Mbytes

Total Flows: 0

Missed Flows: 0

Router	Source	Destination	Source AS	Source AS Name	Dest AS	Dest AS Name

7:16:10AM Done scanning data

By placing the mouse pointer on either the left or the right slider bar and holding down the left mouse button, you can drag the bar to any desired position in the time continuum.

The effect of shortening the applicable time interval is to limit the breadth and scope of FlowAnalyzer operations, which includes reducing the volume of NetFlow data to be processed in responding to user commands.

Conversely, lengthening the applicable time interval for FlowAnalyzer tasks has opposite effects, such as increasing the volume of NetFlow data to be processed, placing heavier demands on system resources, and reducing system performance.

You can position these time slider marks through a wide range to suit the needs of any given FlowAnalyzer task at hand.

- **Get TopN:** pull-down menu—By means of this facility located in the top center of the Display module window (see Figure 1-3), you can select the number of traffic flows that you want taken into account in processing NetFlow data for display purposes.

  For example, you can select the top "N" packets, bytes, or flows for display purposes, where "N" can be any one of the following values: 10, 100, 500, 1000, 2000, 5000, or 10000. The default value for "N" is 100.

- **Sorted by:** pull-down menu—By means of this facility, also located in the top center of the Display module window (see Figure 1-3), you can determine the sort key by which NetFlow data is to be sorted during data retrieval operations. This pull-down menu enables you to "filter" the NetFlow data as octets, packets, or flows for display purposes.
- Status bar—By means of this messaging and status reporting facility at the bottom of the Display module window (Figure 1-3), you can see at a glance the status of any in-process or completed FlowAnalyzer task.

Other user-oriented pull-down menus and buttons incorporated into the Display module window will be covered in Chapter 3, which describes how to use the FlowAnalyzer for a wide variety of NetFlow data retrieval and analysis tasks.

DisplayServer Module

A DisplayServer module running on a host workstation in your network receives and acts on requests for Netflow data that you issue at the console of a Display module. You can configure any number of DisplayServer modules into your network, as appropriate, to act on user requests for NetFlow data.

The DisplayServer module responds to such user requests by accessing the NetFlow data files stored on a designated FlowCollector workstation in the network and transmitting the requested data to the Display module for presentation on the screen of a host workstation or PC. The Display module formats the data on the screen according to the selected aggregation scheme.

UtilityServer Module

The UtilityServer module provides host and autonomous system (AS), as well as device interface information to the FlowAnalyzer.

TrawhoisServer Module

The TrawhoisServer module receives requests from the UtilityServer module to perform AS translation tasks. In response, the TrawhoisServer module returns AS names (as contained in the router arbiter database) to the UtilityServer module.

"Trawhois" is an acronym for trivial routing arbiter whois.