



Analyzer Data Exporting and Collecting Entities

This appendix describes the network traffic data exporting and collection entities that work in combination with the Analyzer to form an integrated network management suite. These network management tools enable network managers, planners, and troubleshooters to better manage and fine tune complex networks.

This appendix describes the following data exporting and collection tools:

- NetFlow DataExporter Devices, page B-1
- TMS DataExporter Devices, page B-4
- NetFlow FlowCollector Hosts, page B-6

NetFlow DataExporter Devices

A NetFlow export-enabled device is one that has been configured to operate with Cisco IOS NetFlow Services software (see Appendix A) in a way that enables the device to export information about traffic flows between communicating end nodes in a network.

A traffic flow is defined as a unidirectional sequence of packets being transmitted between a source node and a destination node in a network.

For NetFlow data export, traffic flows in a network have the following attributes in common:

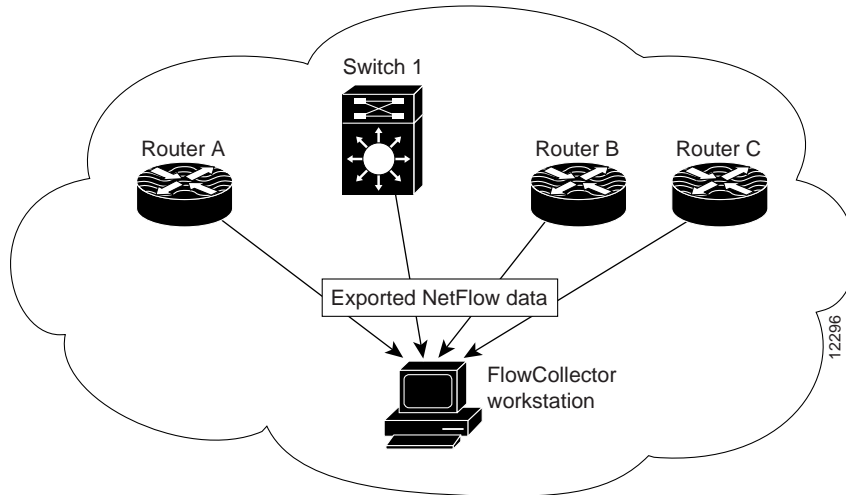
- Source and destination autonomous system (AS) numbers
- Source and destination IP addresses
- Source and destination application port numbers
- Input interfaces
- IP type of services (ToS)
- IP protocol

You can configure any number of NetFlow DataExporter devices into your network to operate in conjunction with Cisco IOS NetFlow Services software (see Appendix A). Such devices enable you to capture and export NetFlow traffic information for later retrieval and analysis.

Figure 1-2 and Figure 1-3 show how such NetFlow DataExporter devices work in combination with other components, enabling you to capture, store, and analyze NetFlow traffic data.

Figure B-1 indicates that multiple NetFlow Data Exporter devices can operate in your network to export NetFlow data to one or more FlowCollector hosts. The FlowCollector passively listens to the UDP ports of the export-enabled devices to collect the NetFlow traffic data periodically exported from the devices.

Figure B-1 NetFlow Export-Enabled Devices in a Network



Network Topology for Capturing NetFlow Data

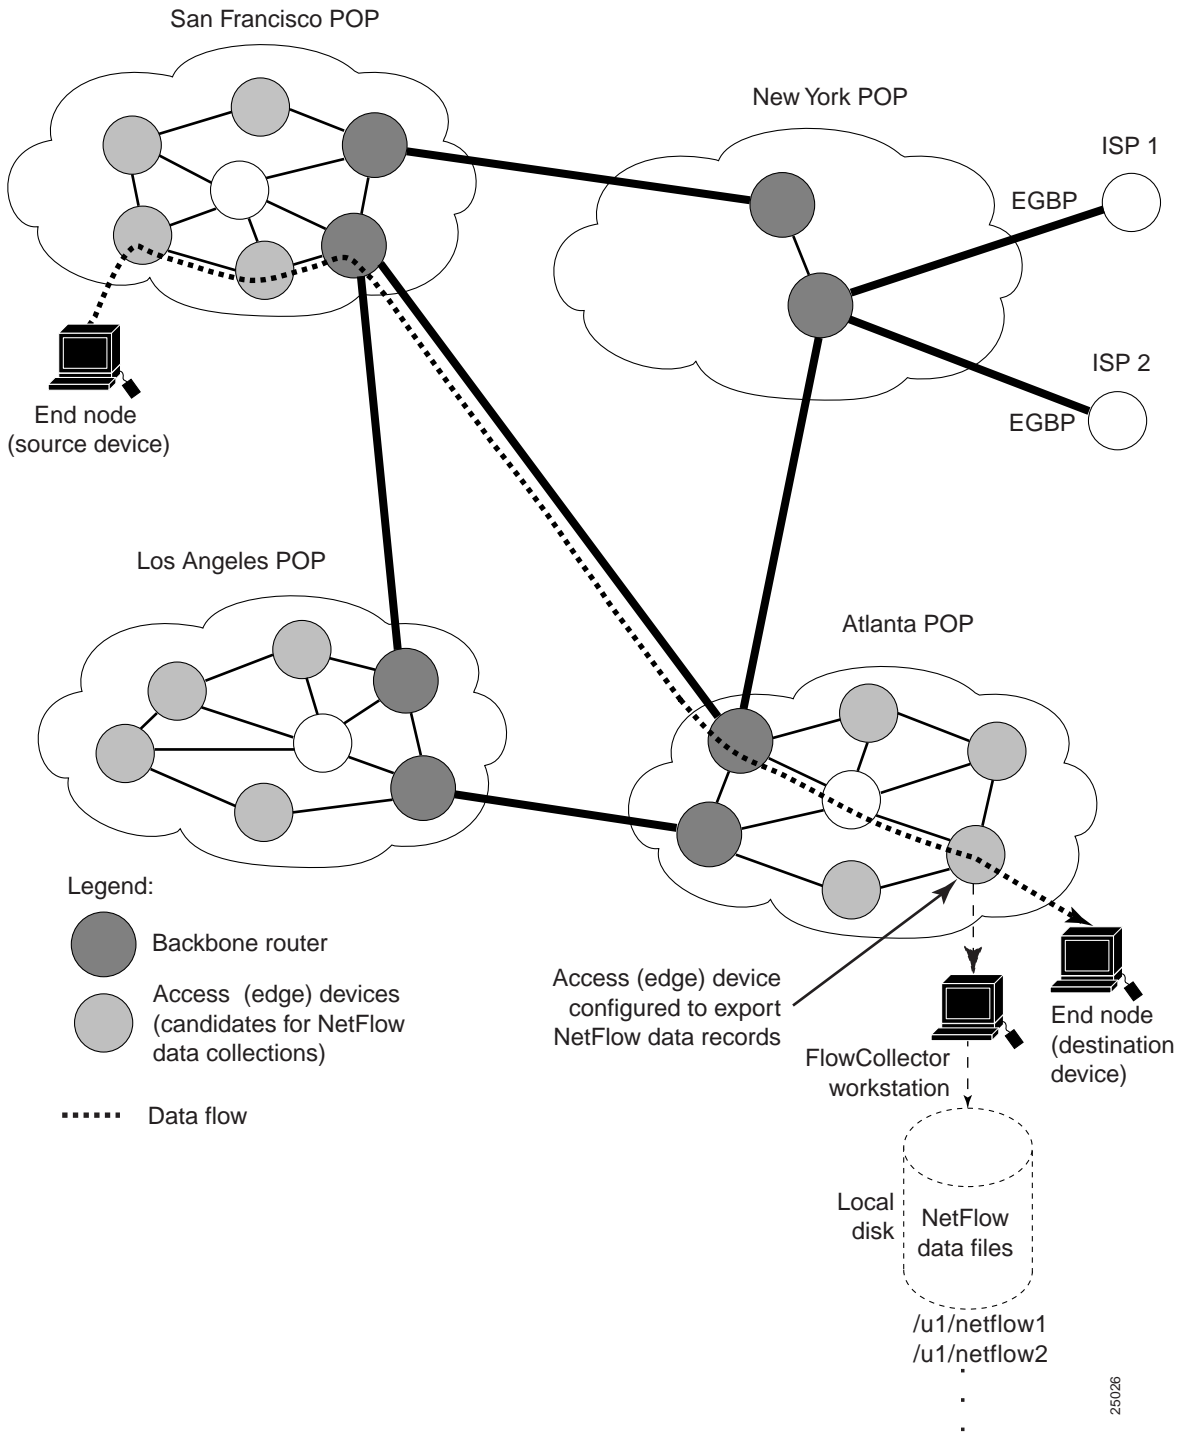
A typical network topology for capturing NetFlow data is shown in Figure B-2. This illustration shows traffic passing between two widely separated end nodes in a complex network. Any NetFlow export-capable device along this or any other transmission path in the network can be configured to capture traffic information pertaining to the communicating devices in an internetworking environment.

Note that the access router (edge device) at the edge of the Atlanta POP has been configured to operate with Cisco IOS NetFlow Services software, thus making it a “NetFlow export-enabled device.” Such a device is capable of monitoring and evaluating the network traffic flowing between the communicating source and destination nodes in the sample network.

As traffic passes through this access router, information from the first packet in the traffic flow is used to build an entry in the NetFlow cache of the export-enabled device. Subsequent packets in the flow are examined by another task that concurrently handles network switching functions, Cisco IOS NetFlow Services functions, and NetFlow data export functions.

NetFlow export records are created according to the Cisco IOS NetFlow Services evaluation criteria, temporarily held in the device’s NetFlow cache, and then exported to the FlowCollector when export criteria are satisfied.

Figure B-2 Representative Network Topology for Capturing NetFlow Data



NetFlow Cache Management

A highly intelligent NetFlow cache management algorithm is employed to ensure the scalability and performance of NetFlow switching functions. These attributes are especially important in the case of densely populated and busy access (edge) routers that handle large numbers of concurrent, short-duration traffic flows.

NetFlow cache management functions include the following:

- Determining if a packet is part of an existing traffic flow, or if the packet should generate a new NetFlow cache entry
- Dynamically updating the per-flow statistics in the NetFlow cache
- Determining the aging, idleness, and storage limits of flows

Long-lived flows, flows that are idle for a specified time period, and flows that exceed the storage limits of the NetFlow cache are aged and removed from the cache.

Based on these cache management functions, Cisco IOS NetFlow Services software assembles the data regarding traffic flows into NetFlow UDP datagrams for export to a FlowCollector host.

Typically, UDP datagrams are exported at least once per second, or whenever a complete NetFlow data export record has been assembled.

NetFlow data export functionality is configured on a per interface basis. For example, to configure a network device to export NetFlow data records, you need only specify the IP address and the application port number of the target FlowCollector host.

TMS DataExporter Devices

You can also configure any number of TMS DataExporter devices (backbone routers) in your network to operate in conjunction with Cisco IOS software. Backbone routers are typically located in the core of a network to aggregate traffic from access devices (edge routers and/or switches) in the network.

Figure 1-2 and Figure 1-3 show how such TMS DataExporter devices work in combination with other components, enabling you to capture, store, and analyze TMS traffic data.

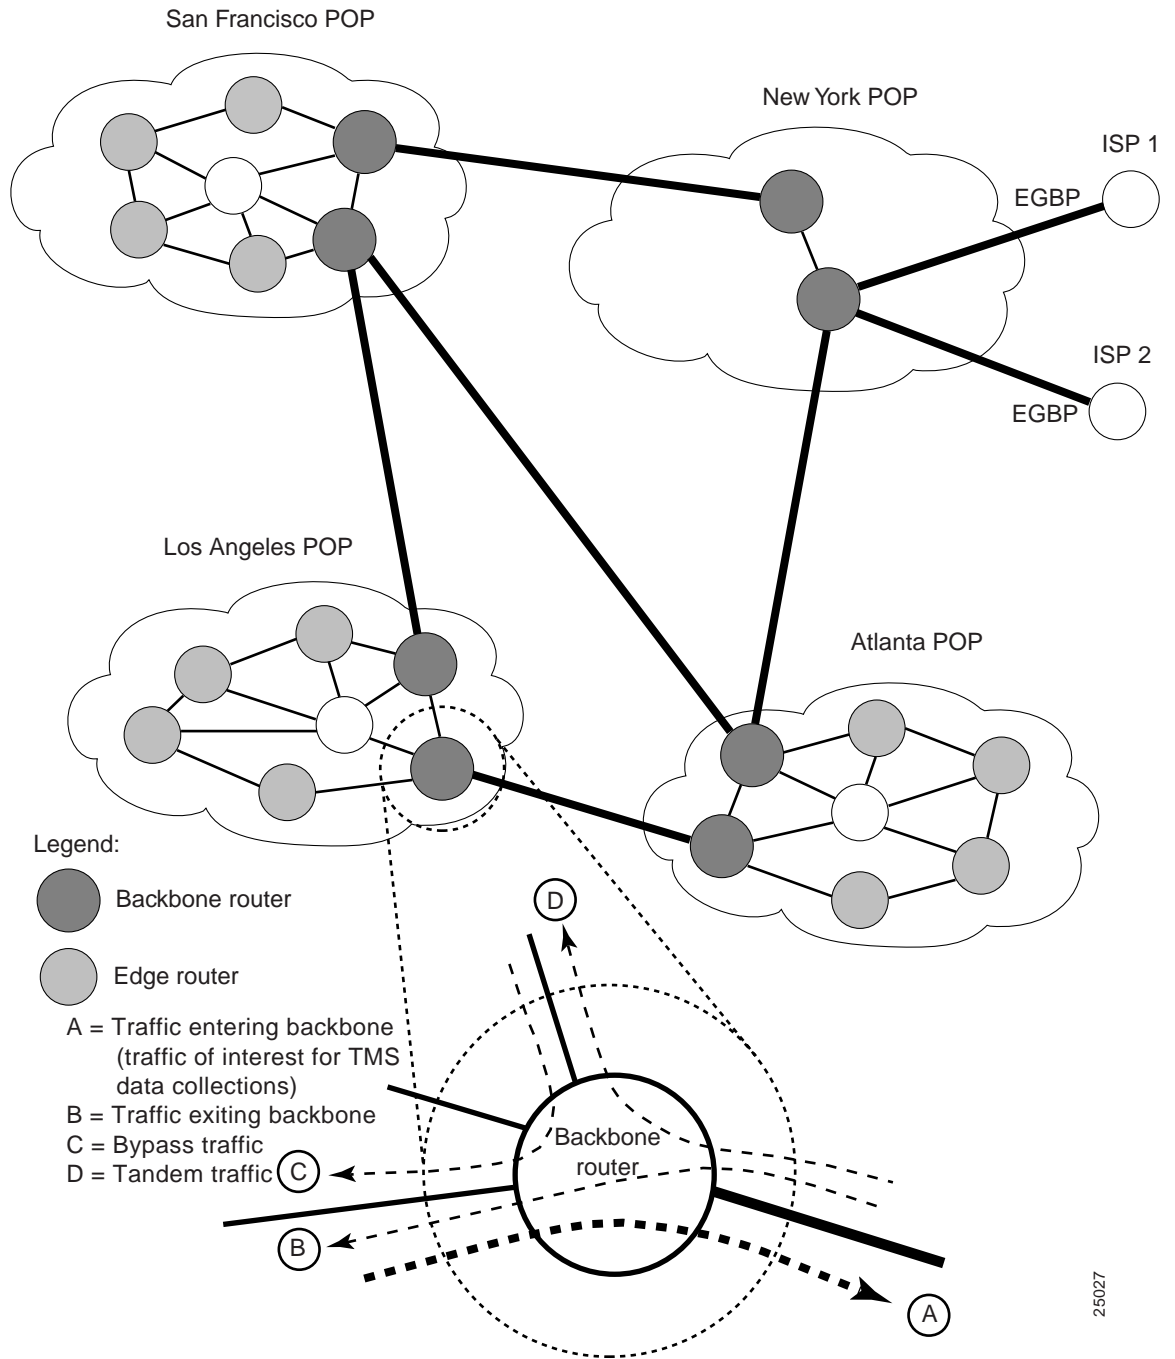
A typical network topology for capturing TMS data is shown in Figure B-3.

TMS data exporting devices periodically capture “snapshots” of TMS traffic information. Free-running counters in the TMS exporting devices are updated dynamically as network traffic passes through the backbone routers.

Periodically, on a user-determined schedule, the UtilityServer is directed to “wake up” and issue a data export command to a specific device. In response, the target export device sends a “snapshot” of its free-running counters to an NFS-mounted storage volume in the network.

Then, when you issue a TMS data display request at the Display module console, the DisplayServer fields the request, retrieves the desired data from the appropriate TMS storage directory, and passes the data on to the Display module for presentation on the console screen in the selected aggregation scheme format.

Figure B-3 Representative Network Topology for TMS Data Collections



The exploded view of the backbone router in Figure B-3 shows the paths of traffic flowing through the device. The relevant path for TMS data is the one introducing traffic onto the trunk line from an external interface (that is, the path labeled "A").

The potential value from collecting traffic data about such a path is that it enables you to determine the relative volume of data being introduced onto the trunk line from an external device.

In a TMS data collection, as traffic passes through a TMS export-enabled backbone router, counters are continuously being incremented to reflect packet flow. Based on parameters specified for the TMS collection, the UtilityServer module (which functions as the schedule keeper for TMS data collections) “wakes up” and issues a command to the backbone router to export data to the appropriate remote storage device. This process is repeated on a periodic, user-determined basis.

Thus, “snapshots” of counter values in the backbone router are captured periodically and exported to the designated storage facility.

NetFlow FlowCollector Hosts

The NetFlow FlowCollector, a client/server application that runs on Solaris platforms, supports fast, scalable, and economical collection of traffic data from one or more NetFlow export-enabled devices in your network.

The FlowCollector, which can be configured to run on one or more workstations, supports the following functionality:

- Collecting NetFlow data from any number of NetFlow export-enabled devices in the network
- Reducing the volume of the collected data by means of selected filters and aggregation schemes
- Storing the collected data in a user-defined directory on a FlowCollector host
- Managing the disk storage space on the FlowCollector host

Once the NetFlow data exporting devices in your network have been configured to operate as desired, the NetFlow FlowCollector passively listens to specified UDP ports to receive the UDP datagrams being exported periodically from the export-enabled devices in your network.

The following sections provide a brief summary of FlowCollector functions and capabilities.

For a comprehensive description of the NetFlow FlowCollector application, see the *NetFlow FlowCollector Installation and User Guide*.

NetFlow Traffic Statistics

NetFlow data records contain detailed traffic information pertaining to traffic flows between communicating end nodes in a network. Such information includes statistics about Layer 3 source and destination nodes, down to the level of the application port numbers and the protocols used by the communicating end nodes.

The ability to collect, store, display, and analyze NetFlow traffic provides the following user benefits:

- Enables network managers to monitor network traffic, determine bandwidth requirements, ensure quality of service (QoS) compliance, and fine-tune network performance
- Enables traffic information to be consolidated and used for billing purposes on a per-application or usage basis

Network traffic statistics typically include the following kinds of information:

- Time stamp of the data flow
- Source and destination IP addresses
- Source and destination port numbers
- Application port numbers
- Next hop addresses

- Total number of bytes, packets, or octets in a flow
- First and last time stamps of packets switched as part of a flow
- Sequence numbers
- Source and destination autonomous system (AS) numbers
- Source and destination prefix masks

NetFlow Traffic Filters

To customize traffic statistics for later display and analysis, you can apply the following filters to the traffic data being collected and stored by the FlowCollector:

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

NetFlow Data Aggregation Schemes

The FlowCollector receives the UDP export datagrams from the NetFlow DataExporter devices in your network, filters the information as directed, and stores the data in a user-specified directory on a FlowCollector host in the format of the specified aggregation scheme.

The NetFlow data aggregation schemes used in conjunction with the FlowCollector are described in the section entitled “NetFlow Data Aggregation Schemes” in Chapter 3.

Data Storage and Management Utilities

The FlowCollector incorporates a powerful set of data storage and management utilities that enable the application to do the following:

- Receive the UDP export datagrams from multiple NetFlow export-enabled devices
- Reduce the volume of received data through the selective application of filters and data aggregation schemes
- Store the processed data in a local directory (where it is available to client applications, such as the Analyzer)
- Perform file cleanup and storage reclamation tasks to minimize the consumption of disk storage space

Also, the FlowCollector incorporates numerous predefined filters that you can apply selectively to

- Include certain traffic information in the FlowCollector database
- Exclude certain traffic information from the FlowCollector database