



Traffic Matrix Statistics

Traffic matrix statistics (TMS) is an IOS feature that enables an administrator to capture and analyze traffic data entering a backbone that is running the Border Gateway Protocol (BGP). This feature also allows an administrator to determine the neighbor autonomous systems of a BGP destination.

This document contains the following major sections:

- Overview, page 1
- Supported Platforms, page 10
- Supported Standards, MIBs, and RFCs, page 10
- Prerequisites, page 11
- Configuration Tasks, page 11
- Configuration Examples, page 14
- Glossary, page 15

Overview

The traffic matrix statistics feature allows an administrator to gather the following data:

- The number of packets and bytes that travel across the backbone from internal and external sources. The packets and bytes are called traffic matrix statistics and are useful for determining how much traffic a backbone handles. You can analyze the traffic matrix statistics, using the following methods:
 - Collecting and viewing the TMS data through the Network Data Analyzer's applications.
 - Reading the TMS data that resides on the backbone router.

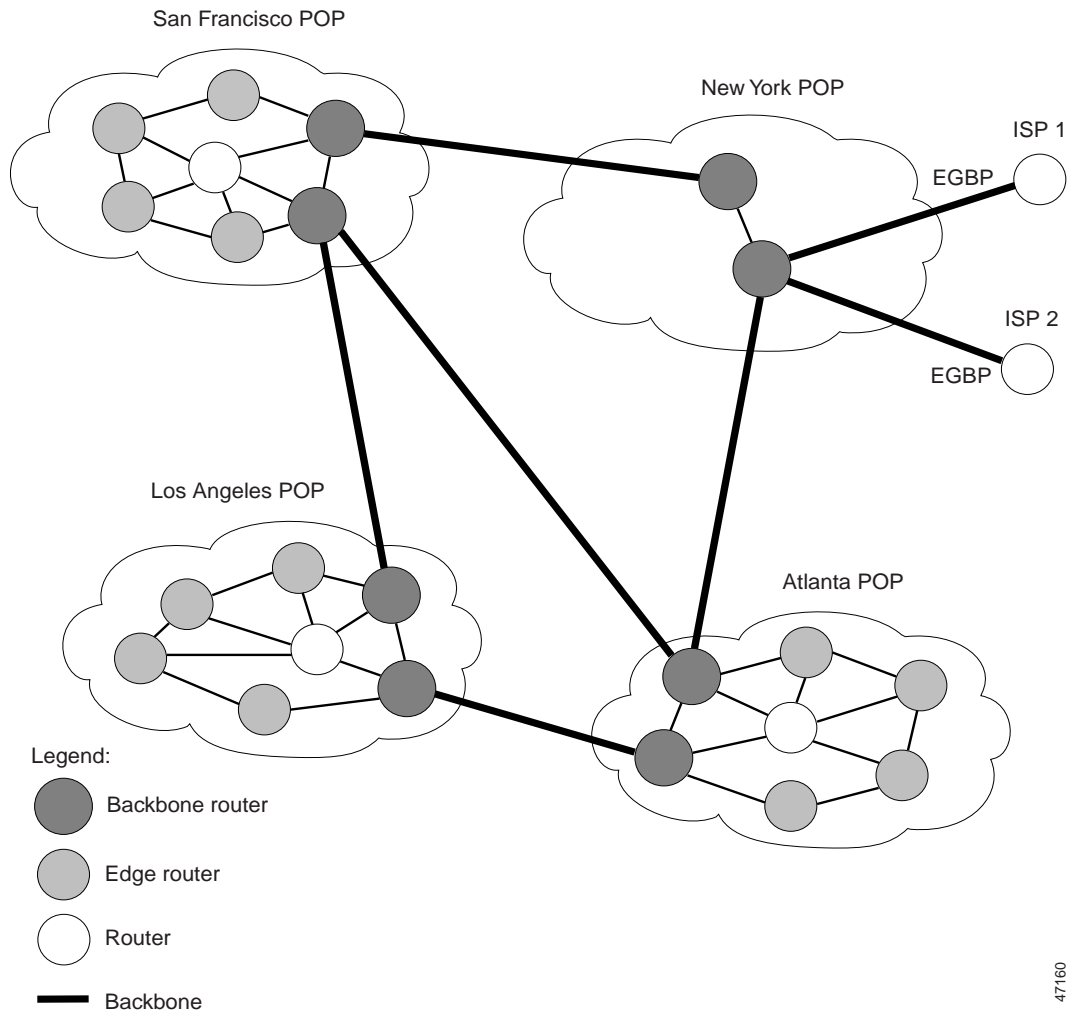
The following sections explain how to collect and view the traffic matrix statistics using the command line interface and the Network Data Analyzer. For detailed instructions on using the Network Data Analyzer, see the *Network Data Analyzer Installation and User Guide*.

- The neighbor autonomous systems of a BGP destination. You can view the neighbor autonomous systems of a BGP destination by reading the `tmasinfo_ascii` file that resides on the backbone router. See the section called Viewing the BGP Neighbor Autonomous Systems, page 9 for more information.

How Backbone Routers Collect TMS Data

By enabling a backbone router to gather traffic matrix statistics, you can determine the amount of traffic that enters the backbone from sites outside of the backbone. You can also determine the amount of traffic that is generated within the backbone. The traffic matrix statistics help you optimize and manage traffic across the backbone. Figure 1 shows a sample backbone, represented by darkly shaded routers and bolded links. The lighter shaded and unshaded routers are outside the backbone. The traffic that travels through the backbone is the area of interest for TMS collection.


Figure 1 Network Backbone and Routers



47160

Figure 2 shows an exploded view of the backbone router that links the Los Angeles point of presence (POP) in Figure 1 to the Atlanta POP. The bold line represents the backbone link going to the Atlanta POP.

Figure 2 Traffic Traveling Through a Backbone Router



The following types of traffic travel through the backbone router shown in Figure 2:

- The dotted line marked A represents traffic entering the backbone from a router that is not part of the backbone. This is called external traffic.
- The dotted lines marked B and D represent traffic that is exiting the backbone. The router interprets traffic from paths B and D as being generated from within the backbone. This is called internal traffic.
- The dotted line marked C represents traffic that is not using the backbone and is not of interest to TMS.

You can determine the amount of traffic the backbone handles by enabling a backbone router to track the number of packets and bytes that travel through it. You can separate the traffic into the categories “internal” and “external.” You separate the traffic by designating incoming interfaces on the backbone router as internal or external.

Once you enable a backbone router to collect traffic matrix statistics, it starts free running counters, which dynamically update when network traffic passes through the backbone router. You can retrieve a snapshot of the traffic matrix statistics, either through a command to the backbone router or through the Network Data Analyzer.

External traffic (Path A) is the most important for determining the amount of traffic. Internal traffic (paths B and D) is useful for ensuring that you are capturing all the TMS data. When you receive a snapshot of the traffic matrix statistics, the packets and bytes are displayed in internal and external categories.

For information on specifying which interfaces on the backbone router should use to collect internal or external traffic, see the section *Enabling a Backbone Router to Collect TMS Data*, page 11.

TMS and CEF Nonrecursive Accounting

TMS data is counted during packet forwarding by CEF nonrecursive accounting, which is configured as described in the Configuration Tasks section. The following paragraphs explain how CEF nonrecursive accounting aggregates packet statistics for IGP routes and their dependent BGP routes.

For example, a BGP network deployed by a service provider has the following components:

- IGP routes that describe the next hop to which traffic should be sent.
- BGP routes that specify an intermediate address to which traffic should be sent.

In this example, the intermediate address might be several hops away. The next hop for the BGP route is the next hop for the BGP route's intermediate address. The BGP route is called recursive, because it points (through its intermediate address) to an IGP route that provides the next hop for forwarding.

CEF represents IGP routes as nonrecursive entries and BGP routes as recursive entries that resolve to nonrecursive entries.

CEF nonrecursive accounting counts the packets for all the CEF recursive entries that resolve to a CEF non-recursive entry and the packets for the non-recursive entry. The number of packets is collected and totalled in one location.

The following example shows how CEF nonrecursive accounting counts packets when BGP routes resolve to one IGP route and when they do not. A multiaccess network access point (NAP) has BGP routes referring to hosts on that network.

- If the network is advertised as a single IGP route, all the BGP routes to the various hosts at that NAP resolve to a single IGP route. CEF nonrecursive accounting summarizes the packets to all of the BGP destinations.
- If a network administrator instead advertises individual host routes from the NAP network to the IGP, CEF nonrecursive accounting will count packets to those hosts separately.

The count of packets forwarded based on a nonrecursive CEF entry can be split into two bins based on whether the backbone router's input interface is configured as internal or external. Thus, all packets that arrive on external interfaces (external to the region of interest) and are forwarded based on a given IGP route (either directly or through a recursive BGP route) are counted together.

For information on configuring a router's input interface as internal or external to a region being analyzed are described in the Configuration Tasks section.

Viewing the TMS Data

Once TMS data is collected, you have the following options for viewing the data:

- Viewing the data in a graphical format, using the Network Data Analyzer Display module. The Display module is useful for graphing the traffic matrix data and comparing statistics. See the section Viewing the TMS Data Through the Network Data Analyzer for more information.
- Issuing the **more system:vfiles/tmstats_ascii** command on the backbone router. This command displays a table of traffic matrix statistics. See the section Viewing the TMS Data by Reading the Virtual Files that Reside on the Backbone Router for more information.
- Issuing the **show ip cef** command on the backbone router. This command displays nonrecursive accounting data for the backbone router. Included in the output is the number of packets and bytes of internal and external traffic have been collected. See the section Viewing TMS Data Through the show ip cef Command for more information.

Viewing the TMS Data Through the Network Data Analyzer

The Network Data Analyzer collects TMS data from the backbone router and displays it using the NDA Display module. The TMS data can look similar to the data shown in Figure 3 and Figure 4. The display format depends on the aggregation scheme you selected. See the *Network Data Analyzer Installation and User Guide* for more information.

(The Network Data Analyzer's Display module is wide. You must slide the scroll bar to the right and left to see all of the data. Figure 3 and Figure 4 taken together show all the columns of data.)

Figure 3 Displaying TMS Data through the Network Data Analyzer (Part 1)

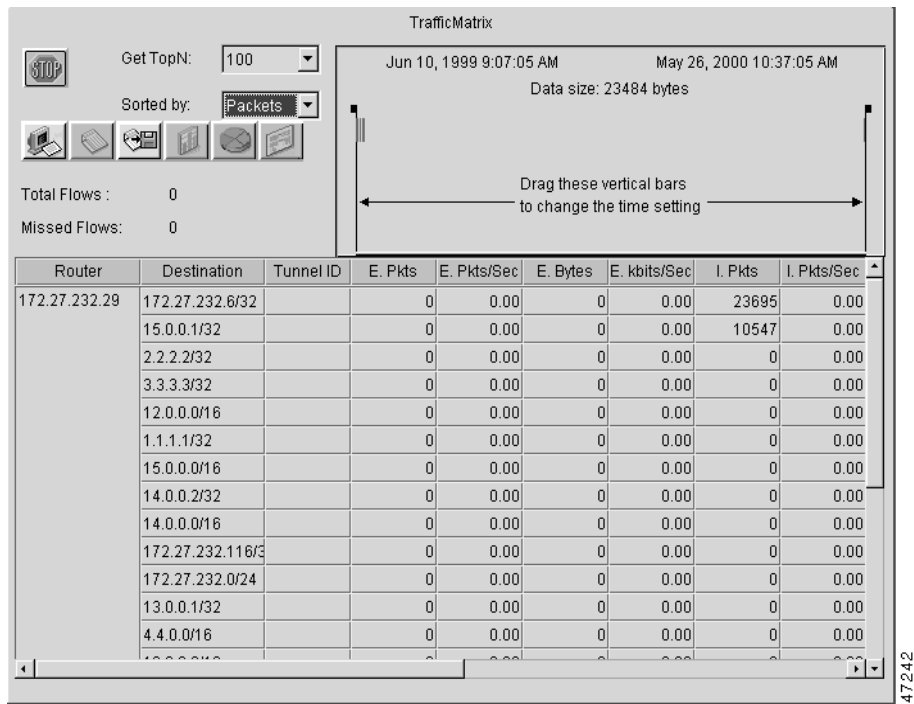
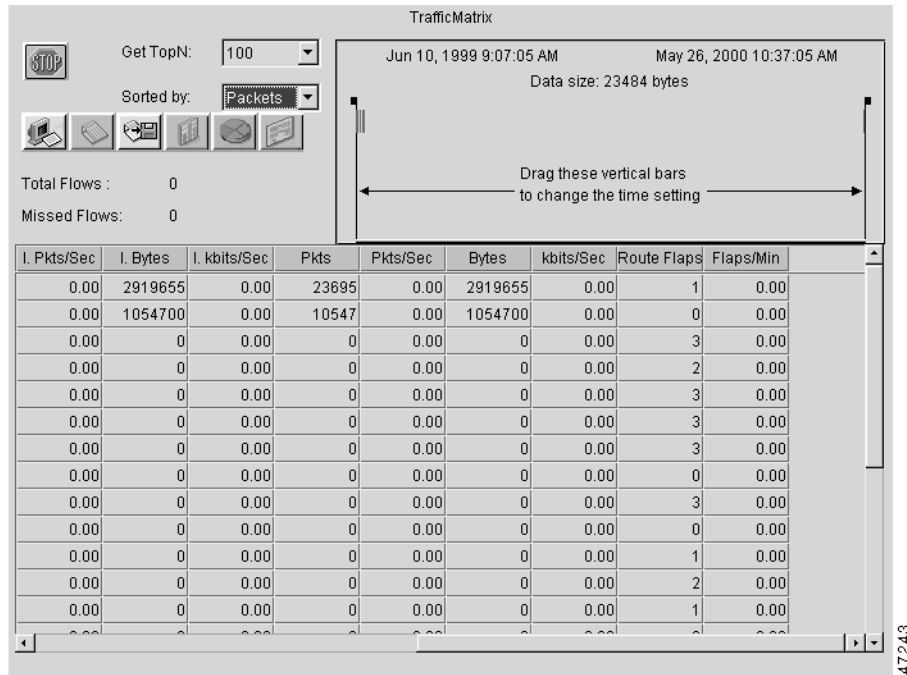


Figure 4 Displaying TMS Data through the Network Data Analyzer (Part 2)



Viewing the TMS Data by Reading the Virtual Files that Reside on the Backbone Router

You can read the TMS data that resides on the backbone router and is stored in the following virtual files:

tmstats_ascii TMS data in ASCII (human readable) format.

tmstats_binary TMS data in binary (space-efficient) format.

Reading the ASCII File

To view statistics in the ASCII file, issue the following command on the backbone router:

```
Router# more system:/vfiles/tmstats_ascii
```

Each file displayed consists of header information and records. A line of space follows the header and each record. A bar (|) separates consecutive fields within a header or record. The first field in a record specifies the type of record. The following example shows a sample TMSTATS_ASCII file:

```
VERSION 1|ADDR 172.27.32.24|AGGREGATION TrafficMatrix.ascii|SYSUPTIME 41428|routerUTC
3104467160|NTP unsynchronized|DURATION 1|
p|10.1.0.0/16|242|1|50|2|100
p|172.27.32.0/22|242|0|0|0|0
```

The following sections describe the header and the various types of records you can display.

File Header

The ASCII file header provides the backbone router's address and information about how much time the router used to collect and export the TMS data. The header occupies one line and uses the following format:

```
VERSION 1|ADDR <address>|AGGREGATION TrafficMatrix.ascii|SYSUPTIME <seconds>|routerUTC
<routerUTC>|NTP <synchronized|unsynchronized>|DURATION <aggregateTime>|
```

Table 1 describes the fields in the file header of the TMSTATS_ASCII file.

Table 1 TMSTATS_ASCII File Header

Maximum Field Length	Field	Description
10	VERSION	File format version
21	ADDR	The IP address of the router
32	AGGREGATION	The type of data being aggregated
21	SYSUPTIME	The time of export, in seconds since the router booted
21	routerUTC	The time of export, in seconds since 1900-01-01 (Coordinated Universal Time (UTC)), as determined by the router
19	NTP	Whether router's Coordinated Universal Time (UTC) has been synchronized by the Network Time Protocol (NTP).
20	DURATION	The time needed to capture the data, in seconds

Destination Prefix Record

The destination prefix record displays the internal and external packets and bytes for the IGP route and uses the following format:

```
p|<destPrefix/Mask>|<creationSysUpTime>|
<internalPackets>|<internalBytes>|<externalPackets>|<externalBytes>
```

Table 2 describes the fields in the destination prefix record.

Table 2 Destination Prefix Record Fields

Maximum Field Length	Field	Description
2	recordType	p means that the record represents dynamic label switching data or traffic engineered (TE) tunnel traffic data.
19	destPrefix/Mask	The IP prefix address/mask (a.b.c.d/len format) for this IGP route.
11	creationSysUpTime	The sysUpTime when the record was first created.
21	internalPackets	Internal packet count.
21	internalBytes	Internal byte count.
21	externalPackets	External packet count.
20	externalBytes	External byte count (no trailing).

Tunnel Midpoint Record

The tunnel midpoint record displays the internal and external packets and bytes for the tunnel head and uses the following format:

```
t|<headAddr> <tun_id>|<creationSysUpTime>|
<internalPackets>|<internalBytes>|<externalPackets>|<externalBytes>
```

Table 3 describes the fields in the tunnel midpoint record.

Table 3 Tunnel Midpoint Record Fields

Maximum Field Length	Field	Description
2	recordType	t means that the record represents traffic engineered (TE) tunnel midpoint data.
27	headAddr<space>tun_id	The IP address of the tunnel head and tunnel interface number.
11	creationSysUpTime	The sysUpTime when the record was first created.
21	internalPackets	Internal packet count.
21	internalBytes	Internal byte count.
21	externalPackets	External packet count.
20	externalBytes	External byte count (no trailing).

Reading the Binary File

The binary file tmstats_binary contains the same information as the ASCII file, except in a space-efficient format. You can copy this file from the router and read it with any utility that accepts files in binary format.

Viewing TMS Data Through the show ip cef Command

You can use the **show ip cef** command to display nonrecursive accounting information, including the internal and external packets and bytes that have traveled through the IP prefix address/mask (a.b.c.d/len format) for an IGP route.

```
router# show ip cef 192.168.1.8
192.168.1.8/32, version 220, per-destination sharing
0 packets, 0 bytes
tag information set
local tag:17
via 192.168.67.8, FastEthernet6/0, 0 dependencies
next hop 192.168.67.8, FastEthernet6/0
valid adjacency
tag rewrite with Fa6/0, 192.168.67.8, tags imposed {}
1143 packets, 56702 bytes switched through the prefix
30 second output rate 0 Kbits/sec
tmstats:external 0 packets, 0 bytes
internal 1144 packets, 56742 bytes
```

Viewing the BGP Neighbor Autonomous Systems

The TMS feature also displays the BGP neighbor autonomous system (AS) associated with each IGP destination. You can display all the neighbor ASs for any IGP destination.

The tmasinfo file is in the ASCII format, which is the only one provided for this data. Issue the following command to read the tmasinfo file:

```
Router# more system:/vfiles/tmasinfo
```

Each file consists of header information and a number of records. A line of space follows the header and each record. A bar (|) separates consecutive fields within a header or a record.

Header Format

The file header provides the router's address and indicates how much time the router used to collect and export the data. The file header uses the following format:

```
VERSION 1|ADDR <address>|AGGREGATION ASList.ascii|SYSUPTIME <seconds>|routerUTC
<routerUTC>|DURATION <aggregateTime>|
```

Table 4 describes the fields in the file header.

Table 4 TMASINFO File Header

Max. Length	Field	Description
5	VERSION	File format version.
15	ADDR	The IP address of the router.
20	AGGREGATION	The type of data being aggregated.
10	SYSUPTIME	The time of export, in seconds since router booted.
10	routerUTC	The time of export, in seconds since 1900-01-01, as determined by router.
10	DURATION	The time needed to capture the data.

Neighbor AS Record

The neighbor AS record displays the neighbor AS and the underlying prefix/mask for each BGP route. The record uses the following format:

```
<nonrecursivePrefix/Mask>|<AS>|<destinationPrefix/Mask>
```

Table 5 describes the fields in the neighbor AS record.

Table 5 Neighbor AS Record Fields

Maximum Field Length	Field	Description
18	nonrecursivePrefix/Mask	The IP prefix address/mask (a.b.c.d/len format) for this IGP route
5	AS	The neighbor AS
18	destinationPrefix/Mask	The prefix/mask for the FIB entry (typically BGP route)

Benefits

The TMS data provides the means to perform the following traffic management tasks:

- Load balancing the network
- Troubleshooting and resolving network problems
- Optimizing network performance
- Planning network expansion
- Planning for failure scenarios, especially at BGP peering points

Restrictions

In order for you to obtain traffic matrix statistics for the backbone, the network must include the following components:

- The routers must be part of a BGP/IGP backbone.
- The supported routers are Cisco 7500, 7200, and GSR.
- TMS is supported by the CEF switching mechanism, for both MPLS labeled packets and unlabeled (IP) packets. TMS does not support process switching, fast switching, or Netflow switching technologies, except as integrated with CEF switching.
 - For CEF switched packets, TMS collects IP-to-IP and IP-to-label packets.
 - For label-switched packets, TMS collects label-to-IP and label-to-label packets. Incoming label packets include those with dynamic labels and TSP tunnel midpoint labels.

Related Documents

For additional information about TMS and the Network Data Analyzer, refer to the following publications:

- *Network Data Analyzer Installation and User Guide*
- *Network Data Analyzer Version 3.0 Release Notes*
- Network Data Analyzer online Help system
- *Cisco IOS Switching Services Configuration Guide*, the section “Cisco Express Forwarding”
- Cisco IOS configuration guides and command reference publications

Supported Platforms

The TMS feature is supported by the Cisco 7500, 7200, and GSR (engine 0 and 1) routers.

Supported Standards, MIBs, and RFCs

No new or modified standards, MIBs, or RFCs are supported by this feature.

Prerequisites

Memory Requirements

Non-recursive accounting statistics are collected in a data structure associated with each non-recursive CEF entry, or approximately each IGP route. The number of IGP routes is typically fewer than the number of BGP routes.

The same data structure is used for load sharing, so non-recursive CEF entries with multiple output routes have this data structure whether nonrecursive accounting is enabled or not. The size of the data structure is approximately 300 bytes.

You can determine the approximate memory requirement in bytes, by using the following formula:

$$(\text{\# of nonrecursive, non-loadshared CEF entries}) * 300 \text{ bytes}$$

Configuration Tasks

The following sections explain two methods of configuring the backbone router to collect TMS data:

- Using the Command Line Interface
- Using the Network Data Analyzer

Enabling a Backbone Router to Collect TMS Data

The procedure for enabling a router to collect TMS data includes enabling nonrecursive accounting and setting the interfaces on the router to collect internal or external traffic matrix statistics. When you set the interfaces on the router, the internal and external settings are used only for TMS collection. The interfaces are set to internal by default.



Note

When setting the interfaces to collect internal and external traffic, make sure you set the incoming interfaces, not the outgoing ones.

You can perform these tasks either through the command line interface or through the Network Data Analyzer. The following sections explain each procedure.

Using the Command Line Interface

Use the following procedure to enable a backbone router to collect TMS data and separate internal and external traffic:

	Command	Purpose
Step 1	Router(config)# ip cef	Enables Cisco express forwarding (CEF) on the router.
Step 2	Router(config)# ip cef accounting non-recursive	Enables nonrecursive accounting on the router.

	Command	Purpose
Step 3	Router(config)# interface <i>type number</i>	Specifies the interface on the backbone router that you intend to configure.
Step 4	Router(config-if)# ip cef accounting non-recursive external or Router(config-if)# ip cef accounting non-recursive internal	Sets the specified incoming interface so that it can collect traffic entering the backbone router from external sources. Sets the specified incoming interface so that it can collect internal traffic in the backbone router.

You can repeat steps 3 and 4 for each incoming interface that you want to configure for TMS.

Using the Network Data Analyzer

Alternatively, you can use the Network Data Analyzer (NDA) to enable TMS data collection and set the incoming interfaces on the backbone router to collect internal and external traffic. The *Network Data Analyzer Installation and User Guide* includes specific instructions.

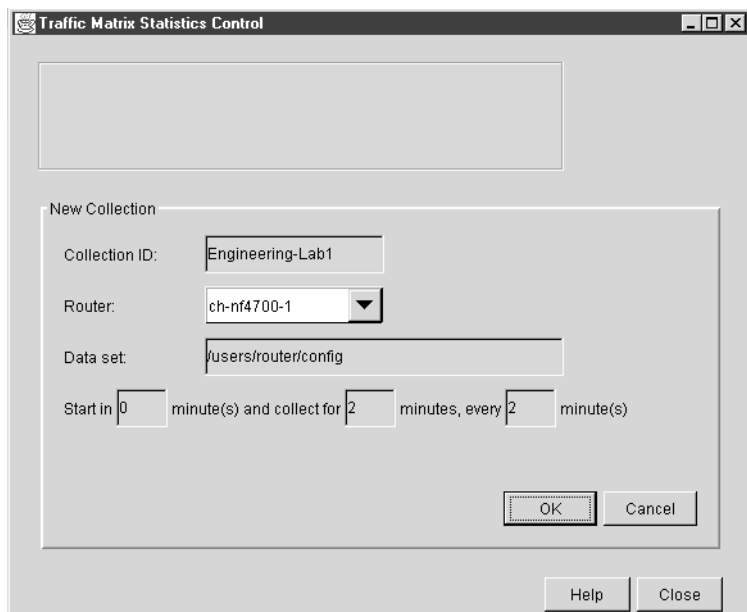
Enabling TMS Data Collection

To enable TMS data collection, you must create a TMS collection and specify:

- The name of the collection
- The router from which you want to collect TMS data
- How often and how long to collect TMS data

The window for enabling the collection of TMS data is similar to the one in Figure 3.

Figure 5 Traffic Matrix Statistics Control Window



Setting Internal and External Interfaces on the Router

The window for setting the interfaces on the backbone router to collect internal and external packets and bytes is similar to the one shown in Figure 6. By default, all interfaces are set to internal. When you set the internal and external interfaces and click Apply, the Network Data Analyzer asks if you want to enable CEF. Click Yes.

Figure 6 Setting Backbone Router Interfaces to Collect Internal and External Traffic Data

The screenshot shows the 'Router Configuration' window for a router named 'ch-nf4700-1' (running IOS version 12.0). The 'NetFlow' tab is selected, and the 'Collecting Data' option is set to 'Yes'. The 'Router Interfaces' section lists several interfaces with radio buttons for 'Internal' and 'External' settings:


Interface	Internal	External
Tunnel2003	<input checked="" type="radio"/>	<input type="radio"/>
Ethernet0	<input type="radio"/>	<input checked="" type="radio"/>
Ethernet1	<input type="radio"/>	<input checked="" type="radio"/>
Ethernet2	<input type="radio"/>	<input checked="" type="radio"/>
Ethernet3	<input type="radio"/>	<input checked="" type="radio"/>
Ethernet4	<input type="radio"/>	<input checked="" type="radio"/>
Ethernet5	<input checked="" type="radio"/>	<input type="radio"/>
FastEthernet0	<input checked="" type="radio"/>	<input type="radio"/>

An 'Apply' button is located in the top right corner of the configuration area. A vertical label '47244' is visible on the right side of the window.

Configuration Examples

This section provides configuration example based on the configuration shown in Figure 7.

Figure 7 Sample Backbone Configuration



The following commands enable the routers to collect internal and external packets and bytes that travel through the backbone routers:

Router A

```

Router(config)# ip cef
Router(config)# ip cef accounting non-recursive
Router(config)# interface e1/0
Router(config-if)# ip cef accounting non-recursive external
  
```

Router B

e1/1:

```

Router(config)# ip cef
Router(config)# ip cef accounting non-recursive
Router(config)# interface e1/1
Router(config-if)# ip cef accounting non-recursive external
  
```

e1/0:

```

Router(config)# interface e1/0
Router(config-if)# ip cef accounting non-recursive internal
  
```

Router C

e1/1:

```

Router(config)# ip cef
Router(config)# ip cef accounting non-recursive
Router(config)# interface e1/1
Router(config-if)# ip cef accounting non-recursive internal
  
```

e1/0:

```

Router(config)# interface e1/0
Router(config-if)# ip cef accounting non-recursive external
  
```

Router D

```

Router(config)# ip cef
Router(config)# ip cef accounting non-recursive
Router(config)# interface e1/1
Router(config-if)# ip cef accounting non-recursive external
  
```

Glossary

Route processor—A general term for a processor module on either the Cisco 7000 or Cisco 7500 router.

Route Processor (RP)—Processor module on a Cisco 7000 series router that contains the CPU, system software and most of the memory components that are used in the router.

Cisco express forwarding (CEF)—CEF is an advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks that have large and dynamic traffic patterns, such as the Internet, as well as for networks characterized by intensive Web-based applications or interactive sessions. CEF uses a Forwarding Information Base (FIB) to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table.

