



# NetFlow Overview

---

**Release 12.1**  
**January 8, 2001**

NetFlow provides network administrators with access to call detail recording information from their data networks. Exported NetFlow data can be used for a variety of purposes, including network management and planning, enterprise accounting and departmental chargebacks, ISP billing, data warehousing and data mining for marketing purposes. NetFlow also provides a highly efficient mechanism with which to process security access lists without paying as much of a performance penalty as is incurred with other available switching methods.

Procedures for configuring NetFlow are provided in the “Configuring NetFlow” chapter later in this publication.

This chapter describes NetFlow. It contains the following sections:

- Accounting Statistics
- NetFlow Data Format
- NetFlow Aggregation
- NetFlow Policy Routing

## Accounting Statistics

NetFlow is a Cisco IOS application that provides statistics on packets flowing through the routing devices in the network. It is emerging as a primary network accounting and security technology.

NetFlow does not involve any connection-setup protocol, either between routers or to any other networking device or end station. NetFlow does not require any change externally—either to the packets themselves or to any networking device. NetFlow is completely transparent to the existing network, including end stations and application software and network devices like LAN switches. Also, NetFlow capture and export are performed independently on each internetworking device; NetFlow need not be operational on each router in the network.

NetFlow is supported on IP and IP encapsulated traffic over most interface types and encapsulations. However, NetFlow does not support ATM LAN emulation (LANE) and does not support an Inter-Switch Link (ISL)/virtual LAN (VLAN), ATM, or Frame Relay interfaces when more than one input access control list (ACL) is used on the interface. Cisco 12000 IP Service Engine ATM line cards do not have this restriction when more than one input ACL is used on the interface.

## Capturing Traffic Data

A network flow is identified as a unidirectional stream of packets between a given source and destination—both defined by a network-layer IP address and transport-layer port number. Specifically, a flow is identified as the combination of the following fields:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Protocol type
- Type of service
- Input interface

## NetFlow Cache

NetFlow operates by creating a flow cache. The NetFlow cache includes entries for traffic statistics that are updated in tandem with the switching of subsequent packets. Information is maintained within the NetFlow cache for all active flows.

## NetFlow Data Format

NetFlow exports flow information in UDP datagrams in one of two formats. The version 1 format was the initially released version, and version 5 is a later enhancement to add Border Gateway Protocol (BGP) autonomous system (AS) information and flow sequence numbers. Versions 2 through 4 were not released.

In version 1 and version 5 formats, the datagram consists of a header and one or more flow records. The first field of the header contains the version number of the export datagram. Typically, a receiving application that accepts either format allocates a buffer big enough for the biggest possible datagram from either format and uses the version from the header to determine how to interpret the datagram. The second field in the header is the number of records in the datagram and should be used to index through the records.

All fields in either version 1 or version 5 formats are in network byte order. Table 5 and Table 6 describe the data format for version 1, and Table 7 and Table 8 describe the data format for version 5.

We recommend that receiving applications check datagrams to ensure that the datagrams are from a valid NetFlow source. We recommend you first check the size of the datagram to make sure it is at least long enough to contain the version and count fields. Next we recommend you verify that the version is valid (1 or 5) and that the number of received bytes is enough for the header and count flow records (using the appropriate version).

Because NetFlow export uses UDP to send export datagrams, it is possible for datagrams to be lost. To determine whether or not flow export information is lost, the version 5 header format contains a flow sequence number. The sequence number is equal to the sequence number of the previous plus the number of flows in the previous datagram. After receiving a new datagram, the receiving application can subtract the expected sequence number from the sequence number in the header to get the number of missed flows.

Table 5 lists the bytes for version 1 header format.

**Table 5** *Version 1 Header Format*

| Bytes | Content           | Description  |
|-------|-------------------|--|
| 0-3   | version and count | Netflow export format version number and number of flows exported in this packet (1-24). |
| 4-7   | SysUptime         | Current time in milliseconds since router booted   |
| 8-11  | unix_secs         | Current seconds since 0000 UTC 1970.   |
| 12-15 | unix_nsecs        | Residual nanoseconds since 0000 UTC 1970.  |

Table 6 lists the byte definitions for version 1 flow record format.

**Table 6** *Version 1 Flow Record Format*

| Bytes | Content               | Description   |
|-------|-----------------------|---|
| 0-3   | srcaddr               | Source IP address.  |
| 4-7   | dstaddr               | Destination IP address.   |
| 8-11  | nexthop               | Next hop router's IP address.   |
| 12-15 | input and output      | Input and output interface's SNMP index.  |
| 16-19 | dPkts                 | Packets in the flow.  |
| 20-23 | dOctets               | Total number of Layer 3 bytes in the flow's packets.                                  |
| 24-27 | First                 | SysUptime at start of flow.   |
| 28-31 | Last                  | SysUptime at the time the last packet of flow was received.                           |
| 32-35 | srcport and dstport   | TCP/UDP source and destination port number or equivalent.                             |
| 36-39 | pad1, prot, and tos   | Unused (zero) byte, IP protocol (for example, 6=TCP, 17=UDP), and IP type-of-service. |
| 40-43 | flags, pad2, and pad3 | Cumulative OR of TCP flags. Pad 2 and pad 3 are unused (zero) byte.                   |
| 44-47 | reserved              | Unused (zero) bytes.  |

Table 7 lists the byte definitions for version 5 header format.

**Table 7** *Version 5 Header Format*

| Bytes | Content           | Description  |
|-------|-------------------|--|
| 0-3   | version and count | Netflow export format version number and number of flows exported in this packet (1-30). |
| 4-7   | SysUptime         | Current time in milliseconds since router booted   |
| 8-11  | unix_secs         | Current seconds since 0000 UTC 1970.   |
| 12-15 | unix_nsecs        | Residual nanoseconds since 0000 UTC 1970.  |
| 16-19 | flow_sequence     | Sequence counter of total flows seen.  |
| 20-23 | reserved          | Unused (zero) bytes.   |

Table 8 lists the byte definitions for version 5 flow record format.

**Table 8** Version 5 Flow Record Format

| Bytes | Content                        | Description   |
|-------|--------------------------------|---|
| 0-3   | srcaddr                        | Source IP address.  |
| 4-7   | dstaddr                        | Destination IP address.   |
| 8-11  | nexthop                        | Next hop router's IP address.   |
| 12-15 | input and output               | Input and output interface's SNMP index.  |
| 16-19 | dPkts                          | Packets in the flow.  |
| 20-23 | dOctets                        | Total number of Layer 3 bytes in the flow's packets.  |
| 24-27 | First                          | SysUptime at start of flow.   |
| 28-31 | Last                           | SysUptime at the time the last packet of flow was received.   |
| 32-35 | srcport and dstport            | TCP/UDP source and destination port number or equivalent.   |
| 36-39 | pad1, tcp_flags, prot, and tos | Unused (zero) byte, Cumulative OR of TCP flags, IP protocol (for example, 6=TCP, 17=UDP), and IP type-of-service. |
| 40-43 | src_as and dst_as              | AS of the source and destination, either origin or peer.  |
| 44-47 | src_mask, dst_mask, and pad2   | Source and destination address prefix mask bits, pad 2 is unused (zero) bytes.                                    |

## NetFlow Aggregation

By maintaining one or more extra flow caches, called *aggregation caches*, the NetFlow Aggregation feature allows limited aggregation of NetFlow data export streams to be done on a router.



### Note

To collect NetFlow version 8 data export records, use NetFlow FlowCollector version 3.0. Version 2.0 and earlier versions do not support version 8 data export record formats.

## Benefits

The NetFlow Aggregation feature provides the following benefits:

- **Reduced Bandwidth Requirement**—NetFlow aggregation caches reduce the bandwidth required between routers and NetFlow management workstations.
- **Reduced NetFlow Workstation Requirements**—NetFlow aggregation caches reduce the number of NetFlow management workstations required.
- **Improved Router Scalability**—NetFlow aggregation caches improve the scalability of high-flow-per-second routers, such as the 7500 series.

## Aggregation Cache Schemes

The aggregation cache schemes are described in the following sections:

- Autonomous System Aggregation Scheme

- Destination Prefix Aggregation Scheme
- Prefix Aggregation Scheme
- Protocol Port Aggregation Scheme
- Source Prefix Aggregation Scheme
- Aggregation Scheme Fields and Key Fields

You can configure each aggregation cache with its individual cache size, cache age timeout parameter, export destination IP address, and export destination UDP port. As data flows expire in the main NetFlow cache, the flows are added to each enabled aggregation cache. Each aggregation cache contains different field combinations that determine which data flows are grouped. The default aggregation cache size is 4096.

Table 9 lists definitions for the data export record terms used in each aggregation scheme.

**Table 9** Data Export Record Terms and Definitions

| Term                              | Definition   |
|-----------------------------------|--|
| Bytes                             | Number of bytes in the aggregated flows.   |
| Destination BGP Autonomous System | Peer or origin autonomous system of the destination prefix (IP address.)   |
| Destination Interface             | SNMP index of the output interface.  |
| Destination Port                  | Destination UDP or TCP port number.  |
| Destination Prefix                | Destination IP address AND'd with the destination prefix mask.   |
| First                             | System uptime when the first packet was switched.  |
| Flows                             | Number of main cache flows that were aggregated.   |
| Last                              | System uptime when the last packet was switched.   |
| Packets                           | Number of packets in the aggregated flows.   |
| PAD                               | Zero field.  |
| Protocol                          | IP protocol byte.  |
| Source BGP Autonomous System      | Peer or origin autonomous system of the source prefix.   |
| Source Interface                  | SNMP index of the input interface.   |
| Source Port                       | Source UDP or TCP port number if applicable.   |
| Source Prefix                     | Source IP address AND'd with the source prefix mask, or the prefix that the source IP address of the aggregated flows belong to. |

## Autonomous System Aggregation Scheme

The autonomous system aggregation scheme provides significant NetFlow export data volume reduction and generates autonomous system-to-autonomous system traffic flow data. The scheme groups data flows with the same source BGP autonomous system, destination BGP autonomous system, input interface, and output interface. See Figure 10.

The aggregated NetFlow data export records report the following:

- Source and destination BGP autonomous system
- Number of packets

- Number of flows summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Output and input interfaces
- Timestamp when the first packet is switched and timestamp when the last packet is switched

**Figure 10 Autonomous System Aggregation Data Export Format**

|    |    |                  |                       |
|----|----|------------------|-----------------------|
| 0  | 0  | Flows            |                       |
| 4  | 4  | Packets          |                       |
| 8  | 8  | Bytes            |                       |
| 12 | 12 | First time stamp |                       |
| 16 | 16 | Last time stamp  |                       |
| 20 | 20 | Source AS        | Destination AS        |
| 24 | 24 | Source interface | Destination interface |

26462

## Destination Prefix Aggregation Scheme

The Destination Prefix aggregation scheme generates data so that you can examine the destinations of network traffic passing through a NetFlow-enabled device. The scheme groups data flows with the same destination prefix, destination prefix mask, destination BGP autonomous system, and output interface. See Figure 11.

The aggregated NetFlow data export records report the following:

- Destination prefix
- Destination prefix mask
- Destination BGP autonomous system
- Number of flows summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Output interface
- Timestamp when the first packet is switched and timestamp when the last packet is switched

**Figure 11 Destination Prefix Aggregation Data Export Record Format**

|    |    |                       |          |
|----|----|-----------------------|----------|
| 0  | 0  | Flows                 |          |
| 4  | 4  | Packets               |          |
| 8  | 8  | Bytes                 |          |
| 12 | 12 | First time stamp      |          |
| 16 | 16 | Last time stamp       |          |
| 20 | 20 | Destination prefix    |          |
| 24 | 24 | Destination mask bits | PAD      |
| 28 | 28 | Destination interface | Reserved |

463

## Prefix Aggregation Scheme

The Prefix aggregation scheme generates data so that you can examine the sources and destinations of network traffic passing through a NetFlow-enabled device. The scheme groups data flows with the same source prefix, destination prefix, source prefix mask, destination prefix mask, source BGP autonomous system, destination BGP autonomous system, input interface, and output interface. See Figure 12.

The aggregated NetFlow data export records report the following:

- Source and destination prefix
- Source and destination prefix mask
- Source and destination BGP autonomous system
- Number of flows summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Input and output interface
- Timestamp when the first packet is switched and timestamp when the last packet is switched

Figure 12 Prefix Aggregation Data Export Record Format

|    |    |                       |                  |                       |
|----|----|-----------------------|------------------|-----------------------|
| 0  | 0  | Flows                 |                  |                       |
| 4  | 4  | Packets               |                  |                       |
| 8  | 8  | Bytes                 |                  |                       |
| 12 | 12 | First time stamp      |                  |                       |
| 16 | 16 | Last time stamp       |                  |                       |
| 20 | 20 | Source prefix         |                  |                       |
| 24 | 24 | Destination prefix    |                  |                       |
| 28 | 28 | Destination mask bits | Source mask bits | Reserved              |
| 32 | 32 | Source AS             |                  | Destination AS        |
| 36 | 36 | Source interface      |                  | Destination interface |

26464

### Protocol Port Aggregation Scheme

The Protocol Port aggregation scheme generates data so that you can examine network usage by traffic type. The scheme groups data flows with the same IP protocol, source port number, and destination port number when applicable. See Figure 13.

The aggregated NetFlow data export records report the following:

- Source and destination port numbers
- IP protocol (where 6=TCP, 17=UDP, etc.)
- Number of flows summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Timestamp when the first packet is switched and timestamp when the last packet is switched

**Figure 13 Protocol Port Aggregation Data Export Record Format**

|    |    |                  |                  |
|----|----|------------------|------------------|
| 0  | 0  | Flows            |                  |
| 4  | 4  | Packets          |                  |
| 8  | 8  | Bytes            |                  |
| 12 | 12 | First time stamp |                  |
| 16 | 16 | Last time stamp  |                  |
| 20 | 20 | Protocol         | Reserved         |
| 24 | 24 | Source port      | Destination port |

26465

### Source Prefix Aggregation Scheme

The Source Prefix aggregation scheme generates data so that you can examine the sources of network traffic passing through a NetFlow-enabled device. The scheme groups data flows with the same source prefix, source prefix mask, source BGP autonomous system, and input interface. See Figure 14.

The aggregated NetFlow data export records report the following:

- Source prefix
- Source prefix mask
- Source BGP autonomous system
- Number of bytes summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Input interface
- Timestamp when the first packet is switched and timestamp when the last packet is switched

Figure 14 Source Prefix Aggregation Data Export Record Format

|    |    |                  |     |           |
|----|----|------------------|-----|-----------|
| 0  | 0  | Flows            |     |           |
| 4  | 4  | Packets          |     |           |
| 8  | 8  | Bytes            |     |           |
| 12 | 12 | First time stamp |     |           |
| 16 | 16 | Last time stamp  |     |           |
| 20 | 20 | Source prefix    |     |           |
| 24 | 24 | Source mask bits | PAD | Source AS |
| 28 | 28 | Source interface |     | Reserved  |

26466

### Aggregation Scheme Fields and Key Fields

To coordinate flow aggregation on your router, determine the fields from which you want to collect data. Table 10 shows which fields are valid for the different aggregation schemes and which fields are part of the keys. Key fields define a unique flow.

Table 10 Aggregation Scheme Data Fields

| Data Fields                       | Aggregation Schemes |                    |        |               |               |
|-----------------------------------|---------------------|--------------------|--------|---------------|---------------|
|                                   | Autonomous System   | Destination Prefix | Prefix | Protocol Port | Source Prefix |
| Source Prefix                     |                     |                    |        |               |               |
| Destination Prefix                |                     |                    |        |               |               |
| Protocol                          |                     |                    |        | *             |               |
| Type of Service Byte              |                     |                    |        |               |               |
| Source Port                       |                     |                    |        | *             |               |
| Destination Port                  |                     |                    |        | *             |               |
| Source Interface                  | *                   |                    | *      |               | *             |
| Destination Interface             | *                   | *                  | *      |               |               |
| OR'd TCP Flags                    |                     |                    |        |               |               |
| Source BGP Autonomous System      | *                   |                    | *      |               | *             |
| Destination BGP Autonomous System | *                   | *                  | *      |               |               |
| Source Prefix Mask                |                     |                    | *      |               | *             |
| Destination Prefix Mask           |                     | *                  | *      |               |               |
| Next Hop IP Address               |                     |                    |        |               |               |
| Source Encap Bytes                |                     |                    |        |               |               |
| Destination Encap Bytes           |                     |                    |        |               |               |
| Source Prefix                     |                     |                    | *      |               | *             |
| Destination Prefix                |                     | *                  | *      |               |               |
| First Timestamp                   | x                   | x                  | x      | x             | x             |
| Last Timestamp                    | x                   | x                  | x      | x             | x             |
| Flows                             | x                   | x                  | x      | x             | x             |
| Packets                           | x                   | x                  | x      | x             | x             |
| Bytes                             | x                   | x                  | x      | x             | x             |

\* = exported key field

x = exported field

## New Version 8 NetFlow Data Export Support

NetFlow exports flow information in UDP datagrams in one of several formats. Version 8, a new data export version, has been added to support data exports from aggregation caches. Version 8 allows for export datagrams to contain a subset of the usual version 5 export data, which is valid for a particular aggregations scheme type.

Figure 15 shows the version 8 header with the version and timestamp information. Table 11 lists definitions for terms used in the version 8 header.

**Figure 15** Version 8 Header Format

|    |                  |           |             |                     |
|----|------------------|-----------|-------------|---------------------|
| 0  | Version          |           | Count       |                     |
| 4  | System uptime    |           |             |                     |
| 8  | UNIX seconds     |           |             |                     |
| 12 | UNIX nanoseconds |           |             |                     |
| 16 | Sequence number  |           |             |                     |
| 20 | Engine type      | Engine ID | Aggregation | Aggregation version |
| 24 | Reserved         |           |             |                     |

26467

**Table 11** Terms and Definitions for Version 8 Headers

| Term                | Definition  |
|---------------------|---|
| Version             | The flow export format version number. In this case, the number is “8.” |
| Count               | The number of export records in the datagram.                           |
| System Uptime       | The number of milliseconds since the router was last booted.            |
| UNIX Seconds        | The number of seconds since 0000 Universal Time Code (UTC) 1970.        |
| UNIX Nanoseconds    | The number of residual nanoseconds since 0000UTC 1970.                  |
| Sequence Number     | Sequence counter of total flows sent for this export stream.            |
| Engine Type         | The type of switching engine. RP=0 and LC=1.                            |
| Engine ID           | The slot number of the NetFlow engine.                                  |
| Aggregation         | The type of aggregation scheme being used.                              |
| Aggregation Version | The aggregation subformat version number. The current value is “2.”     |

# NetFlow Policy Routing

NetFlow policy routing (NPR) integrates policy routing, which enables traffic engineering and traffic classification, with NetFlow services, which provide billing, capacity planning, and monitoring information on real-time traffic flows. IP policy routing now works with Cisco Express Forwarding (CEF), Distributed CEF (dCEF), and NetFlow.

As Quality of Service and traffic engineering become more popular, so does interest in policy routing's ability to selectively set IP precedence and type of service (TOS) bits (based on access lists and packet size), thereby routing packets based on predefined policy. It is important that policy routing work well in large, dynamic routing environments. Hence, distributed support allows customers to leverage their investment in distributed architecture.

Cisco IOS Software Release 12.1 includes three technologies for IP Policy Routing:

**Table 12** Table 12: Technologies for IP Policy Routing

| Technology | Description   |
|------------|---|
| CEF        | Looks at a Forwarding Information Base (FIB) instead of a routing table when switching packets. |
| dCEF       | Addresses the scalability and maintenance problems of a demand caching scheme.                  |
| NetFlow    | Provides accounting, capacity planning, and traffic monitoring capabilities.                    |

NPR leverages these technologies. To configure NetFlow policy routing, see the chapter "Configuring NetFlow" in this publication.

## Benefits

The NetFlow Policy Routing feature provides the following benefits:

- NPR takes advantage of the new switching services. CEF and dCEF can now use policy routing.
- Now that policy routing is integrated into CEF, policy routing can be deployed on a wide scale and on high-speed interfaces.

## Restrictions

Note the following restrictions:

- NPR is only available on Cisco IOS CEF-based platforms.
- Distributed FIB-based policy-routing is only available on platforms that support dCEF and images that support dCEF.
- The **set ip next-hop verify-availability** command is not supported in dCEF because dCEF does not support the Cisco Discovery Protocol (CDP) database.
- The Cisco 12000 platform does not support this feature in Cisco IOS Release 12.1.