NetFlow Services Solutions Guide

Version History

Version Number	Date	Notes
1	5/9/2001	This document was created.
2	7/16/2001	Incorporated editorial comments.
3	10/23/2004	Updated all sections.
4	01/22/2007	Updated information on types of flow switching engines.

The NetFlow Services Solutions Guide contains the following sections:

Introduction

Rapid growth of IP networks has created interest in new business applications and services. These new services have resulted in increases in demand for network bandwidth, performance, and predictable quality of service as well as VoIP, multimedia and security oriented network services. Simultaneously, the need has emerged for measurement technology to support this growth by efficiently providing the information required to record network and application resource utilization. Cisco's IOS NetFlow provides solutions for each of these challenges.

This white paper is an overview of NetFlow benefits and includes technical overview of features, details about the NetFlow cache, export formats and NetFlow operation.

NetFlow Definitions and Benefits

NetFlow traditionally enables several key customer applications including:

- *Network Monitoring*—NetFlow data enables extensive near real time network monitoring capabilities. Flow-based analysis techniques may be utilized to visualize traffic patterns associated with individual routers and switches as well as on a network-wide basis (providing aggregate traffic or application based views) to provide proactive problem detection, efficient troubleshooting, and rapid problem resolution.

- *Application Monitoring and Profiling*—NetFlow data enables network managers to gain a detailed, time-based, view of application usage over the network. This information is used to plan, understand new services, and allocate network and application resources (e.g. Web server sizing and VoIP deployment) to responsively meet customer demands.
- *User Monitoring and Profiling*—NetFlow data enables network engineers to gain detailed understanding of customer/user utilization of network and application resources. This information may then be utilized to efficiently plan and allocate access, backbone and application resources as well as to detect and resolve potential security and policy violations.
- *Network Planning*—NetFlow can be used to capture data over a long period of time producing the opportunity to track and anticipate network growth and plan upgrades to increase the number of routing devices, ports, or higher- bandwidth interfaces. NetFlow services data optimizes network planning including peering, backbone upgrade planning, and routing policy planning. NetFlow helps to minimize the total cost of network operations while maximizing network performance, capacity, and reliability. NetFlow detects unwanted WAN traffic, validates bandwidth and Quality of Service (QOS) and allows the analysis of new network applications. NetFlow will give you valuable information to reduce the cost of operating your network.
- *Security Analysis*—NetFlow identifies and classifies DDOS attacks, viruses and worms in real-time. Changes in network behavior indicate anomalies that are clearly demonstrated in NetFlow data. The data is also a valuable forensic tool to understand and replay the history of security incidents.
- *Accounting/Billing*—NetFlow data provides fine-grained metering (e.g. flow data includes details such as IP addresses, packet and byte counts, timestamps, type-of-service and application ports, etc.) for highly flexible and detailed resource utilization accounting. Service providers may utilize the information for billing based on time-of-day, bandwidth usage, application usage, quality of service, etc. Enterprise customers may utilize the information for departmental charge-back or cost allocation for resource utilization.
- *NetFlow Data Warehousing and Data Mining*—NetFlow data (or derived information) can be warehoused for later retrieval and analysis in support of proactive marketing and customer service programs (e.g. figure out which applications and services are being utilized by internal and external users and target them for improved service, advertising, etc.). In addition, NetFlow data gives Market Researchers access to the "who", "what", "where", and "how long" information relevant to enterprises and service providers.

NetFlow has two key components: (1) the NetFlow cache or data source which stores IP Flow information and (2) the NetFlow export or transport mechanism that sends NetFlow data to a network management collector for data reporting. The Cisco IOS Flexible and extensible export format, NetFlow version 9, is now on the IETF standards track in the IP Information export (IPFIX) working group. The new generic data transport capability within Cisco routers, IPFIX export, can be used to transport any performance information from a router or switch. The main NetFlow focus has always been IP Flow information but this is now changing with Cisco implementation of a generic export transport format that is an innovative IETF standard. New information is being exported using the NetFlow version 9 export format including Layer 2 information, new security detection and identification information, IPv6, Multicast, MPLS, BGP information, and more.

What Is A Flow?

A 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 source and destination port numbers. Specifically, a flow is identified as the combination of the following seven key fields:

- Source IP address

- Destination IP address
- Source port number
- Destination port number
- Layer 3 protocol type
- ToS byte
- Input logical interface (ifIndex)

These seven key fields define a unique flow. If a flow has one different field than another flow, then it is considered a new flow. A flow contains other accounting fields (such as the AS number in the NetFlow export Version 5 flow format) that depend on the version record format that you configure for export. Flows are processed in a NetFlow cache.

NetFlow Cache Management and Data Export

Building a NetFlow Cache

NetFlow operates by creating a NetFlow cache entry that contains the information for all active flows. The NetFlow cache is built by processing the first packet of a flow through the standard switching path. A Flow record is maintained within the NetFlow cache for all active flows. Each flow record in the NetFlow cache contains key fields that can be later used for exporting data to a collection device. Each flow record is created by identifying packets with similar flow characteristics and counting or tracking the packets and bytes per flow. The flow details or cache information is exported to a flow collector server periodically based upon flow timers. The collector contains a history of flow information that was switched within Cisco device. NetFlow is very efficient, the amount of export data being about 1.5% of the switched traffic in the router. NetFlow accounts for every packet (non-sampled mode) and provides a highly condensed and detailed view of all network traffic that entered the router or switch.

The key to NetFlow-enabled switching scalability and performance is highly intelligent flow cache management, especially for densely populated and busy edge routers handling large numbers of concurrent, short duration flows. The NetFlow cache management software contains a highly sophisticated set of algorithms for efficiently determining if a packet is part of an existing flow or should generate a new flow cache entry. The algorithms are also capable of dynamically updating per-flow accounting measurements residing in the NetFlow cache, and cache aging/flow expiration determination.

Rules for expiring NetFlow cache entries include:

- Flows which have been idle for a specified time are expired and removed from the cache
- Long lived flows are expired and removed from the cache (flows are not allowed to live more than 30 minutes by default, the underlying packet conversation remains undisturbed)
- As the cache becomes full a number of heuristics are applied to aggressively age groups of flows simultaneously
- TCP connections which have reached the end of byte stream (FIN) or which have been reset (RST) will be expired.

Expired flows are grouped together into "NetFlow Export" datagrams for export from the NetFlow-enabled device. NetFlow Export datagrams may consist of up to 30 flow records for version 5 or 9 flow export. NetFlow functionality is configured on a per-interface basis. To configure NetFlow Export capabilities, the user simply needs to specify the IP address and application port number of the Cisco

NetFlow or third-party FlowCollector. The FlowCollector is a device that provides NetFlow Export data filtering and aggregation capabilities. Figure 1 shows an example of the NetFlow cache, aggregation cache and timers.

Figure 1 Example of a NetFlow Cache

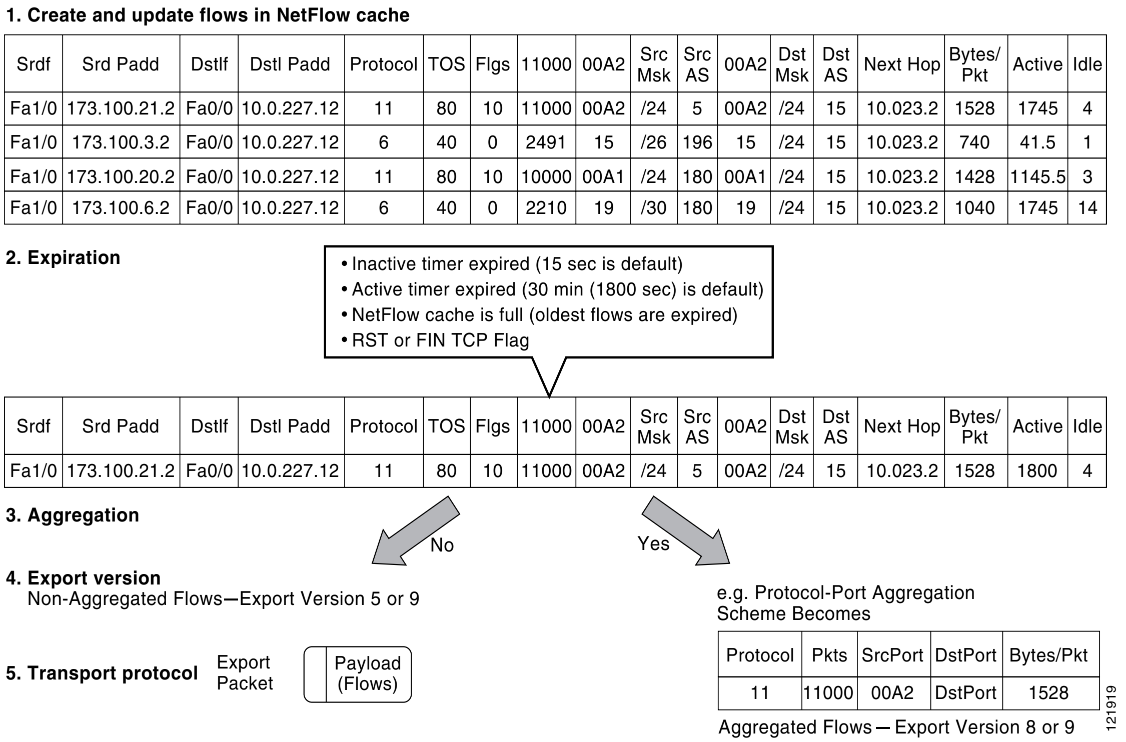

1. Create and update flows in NetFlow cache

Srdf	Srd Padd	DstIf	DstI Padd	Protocol	TOS	Flgs	11000	00A2	Src Msk	Src AS	00A2	Dst Msk	Dst AS	Next Hop	Bytes/ Pkt	Active	Idle
Fa1/0	173.100.21.2	Fa0/0	10.0.227.12	11	80	10	11000	00A2	/24	5	00A2	/24	15	10.023.2	1528	1745	4
Fa1/0	173.100.3.2	Fa0/0	10.0.227.12	6	40	0	2491	15	/26	196	15	/24	15	10.023.2	740	41.5	1
Fa1/0	173.100.20.2	Fa0/0	10.0.227.12	11	80	10	10000	00A1	/24	180	00A1	/24	15	10.023.2	1428	1145.5	3
Fa1/0	173.100.6.2	Fa0/0	10.0.227.12	6	40	0	2210	19	/30	180	19	/24	15	10.023.2	1040	1745	14

2. Expiration

- Inactive timer expired (15 sec is default)
- Active timer expired (30 min (1800 sec) is default)
- NetFlow cache is full (oldest flows are expired)
- RST or FIN TCP Flag

Srdf	Srd Padd	DstIf	DstI Padd	Protocol	TOS	Flgs	11000	00A2	Src Msk	Src AS	00A2	Dst Msk	Dst AS	Next Hop	Bytes/ Pkt	Active	Idle
Fa1/0	173.100.21.2	Fa0/0	10.0.227.12	11	80	10	11000	00A2	/24	5	00A2	/24	15	10.023.2	1528	1800	4

3. Aggregation

4. Export version
Non-Aggregated Flows—Export Version 5 or 9

e.g. Protocol-Port Aggregation
Scheme Becomes

Protocol	Pkts	SrcPort	DstPort	Bytes/Pkt
11	11000	00A2	DstPort	1528

Aggregated Flows—Export Version 8 or 9

5. Transport protocol

NetFlow IOS Packaging Information

Cisco 7200/7500/7400/MGX/AS5800—Although NetFlow functionality is physically included in all software images for these platforms, customers must purchase a separate NetFlow Feature License in order to be licensed for its use. NetFlow licenses are sold on a per-node basis.

Other routers—NetFlow functionality is supported only in Plus images for these platforms. Customers are required to purchase an appropriate Plus image in order to utilize NetFlow functionality on these platforms. There is no feature license for most Cisco platforms except the following require a software license Cisco *7200/7500/7400/MGX/AS5800.*

Reformation IOS Packages—NetFlow is currently available in IP Base package and above.

NetFlow Export Version Formats

The NetFlow Export datagram consists of a header and a sequence of flow records. The header contains information such as sequence number, record count and sysuptime. The flow record contains flow information, for example IP addresses, ports, and routing information. Figure 2 is a typical datagram used for NetFlow fixed format export versions 1, 5, 7 and 8.

Figure 2 Typical NetFlow Export Datagram Format for Versions 1, 5, 7, 8

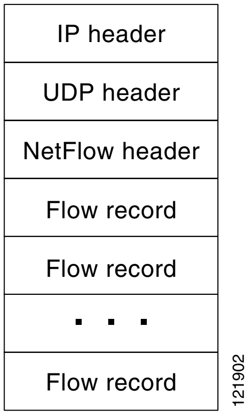

The Version 1 export format was the original format supported in the initial Cisco IOS software releases containing NetFlow functionality and is rarely used today. The Version 5 format is a later enhancement that adds Border Gateway Protocol (BGP) autonomous system information and flow sequence numbers. The Version 7 format is an enhancement that adds NetFlow support for Cisco Catalyst series switches using hybrid or native mode. If you are wondering what happened to Versions 2 through 4 and Version 6 they were either not released or are not supported. Version 8 is the NetFlow export format used when the Router-Based NetFlow Aggregation feature is enabled on Cisco IOS router platforms and is discussed later. The most recent evolution of the NetFlow flow-record format is known as Version 9. The distinguishing feature of the NetFlow Version 9 format is that it is *template based*. Templates provide an extensible design to the record format, a feature that should allow future enhancements to NetFlow services without requiring concurrent changes to the basic flow-record format. NetFlow Version 9 is now the protocol of choice for the IETF IP Information Export (IPFIX) WG and the IETF Pack Sampling WG (PSAMP).

Using templates with NetFlow Version 9 provides several key benefits:

- Almost any information can be exported from a router or switch including layer 2 through 7 information, routing information, IPv6, IPv4, multicast and MPLS information. This new information will allow new applications for flow data and new views of network behavior.
- Third-party business partners who produce applications that provide collector or display services for NetFlow will not be required to recompile their applications each time a new NetFlow export field is added. Instead, they may be able to use an external data file that documents the known template formats.
- New features can be added to NetFlow more quickly, without breaking current implementations.
- NetFlow is "future-proofed" against new or developing protocols, because the Version 9 format can be adapted to provide support for them and other non-Flow based data measurements.

NetFlow Export Packet Header Format

In all five versions, 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 any of the format versions allocates a buffer large enough for the largest possible datagram from any of the format versions and then uses the header to determine how to interpret the datagram. The second field in the header contains the number of records in the datagram (indicating the number of expired flows represented by this datagram) and is used to index through the records. Datagram headers for NetFlow Export versions 5, 7, 8 and 9 also include a "sequence number" field used by NetFlow data consuming applications to check for lost datagrams.

The NetFlow Version 9 export header format is shown below in Figure 3. For additional information see Appendix 1: Details Of NetFlow Export Packet Header Format For Each Export Version.

Figure 3 NetFlow Version 9 Header Format

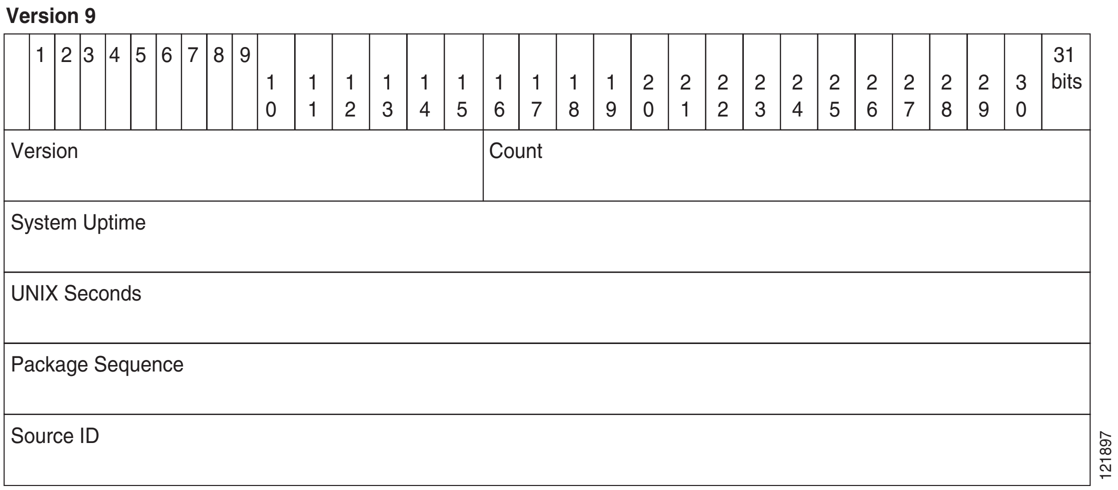

Table 1 shows the field names and values for the Version 9 header format

Table 1 Version 9 Header Format

Field Name	Value
Version	The version of NetFlow records exported in this packet; for Version 9, this value is 0x0009.
Count	Number of FlowSet records (both template and data) contained within this packet.
System Uptime	Time in milliseconds since this device was first booted.
UNIX Seconds	Seconds since 0000 Coordinated Universal Time (UTC) 1970.

Field Name	Value
Sequence Number	Incremental sequence counter of all export packets sent by this export device; this value is cumulative, and it can be used to identify whether any export packets have been missed. This is a change from the NetFlow Version 5 and Version 8 headers, where this number represented "total flows."
Source ID	The Source ID field is a 32-bit value that is used to guarantee uniqueness for all flows exported from a particular device. (The Source ID field is the equivalent of the engine type and engine ID fields found in the NetFlow Version 5 and Version 8 headers. The format of this field is vendor specific. In Cisco's implementation, the first two bytes are reserved for future expansion, and will always be zero. Byte 3 provides uniqueness with respect to the routing engine on the exporting device. Byte 4 provides uniqueness with respect to the particular line card or Versatile Interface Processor on the exporting device. Collector devices should use the combination of the source IP address plus the Source ID field to associate an incoming NetFlow export packet with a unique instance of NetFlow on a particular device.

Cisco IOS NetFlow Flow Record and Export Format Content Information

The section outlines details about the Cisco export format flow record. Table 2 is an example of typical export format fields available for version 5, 7, and 9.

Table 2 Shows the NetFlow flow record contents

Field	Version 5	*Version 5 Catalyst 65k	Version 9	*Version 7 Catalyst 65k
source IP address	Y	Y	Y	Y
destination IP address	Y	Y	Y	Y
source TCP/UDP application port	Y	Y	Y	Y
destination TCP/UDP application port	Y	Y	Y	Y
next hop router IP address	Y	Y 12.1(13)E	Y	Y
input physical interface index	Y	Y	Y	Y
output physical interface index	Y	Y 12.1(13)E	Y	Y
packet count for this flow	Y	Y	Y	Y
byte count for this flow	Y	Y	Y	Y
start of flow timestamp	Y	Y	Y	Y
end of flow timestamp	Y	Y	Y	Y
IP Protocol (for example, TCP=6; UDP=17)	Y	Y	Y	Y
Type of Service (ToS) byte	Y	***PFC3b Only	Y	***PFC3b Only
TCP Flags (cumulative OR of TCP flags)	Y	N	Y	N
source AS number	Y	Y 12.1(13)E	Y	Y 12.1(13)E
destination AS number	Y	Y 12.1(13)E	Y	Y 12.1(13)E

source subnet mask	Y	N	Y	N
destination subnet mask	Y	N	Y	N
flags (indicates, among other things, which flows are invalid)	Y	Y	Y	Y
shortcut router IP address[3]	N	N	N	Y
**Other flow fields	N	N	Y	N

*Assumes use of the full interface flow mask configuration. For more information on fields and flow masks available on the Catalyst 65k see Appendix 2: Details for NetFlow Export Formats.

** For a complete list of other flow fields available in version 9 see NetFlow Version 9 Export Packet Example.

*** TOS is based on first packet in the flow

[3] IP address of the router that is shortcutted by the Catalyst series switch.

Figure 4 is an example of the NetFlow version 5 record format including the contents and description of byte locations.

Figure 4 Exporting the Version 5 Record Format

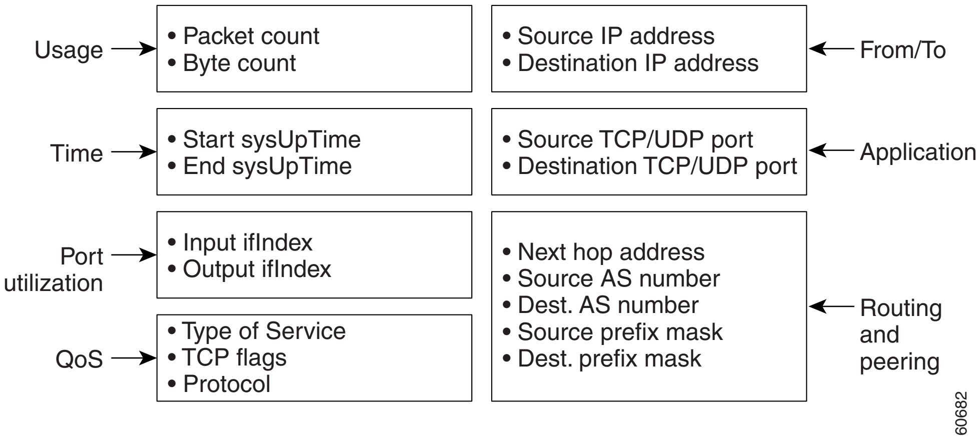

Table 3 Shows the field names and values for the version 5 header format

Table 3 Field Names and Values for the Version 5 Header Format.

Content	**Bytes**	**Description**
srcaddr	0-3	Source IP address
dstaddr	4-7	Destination IP address
nexthop	8-11	Next hop router's IP address
input	12-13	Ingress interface SNMP ifIndex
output	14-15	Egress interface SNMP ifIndex
dPkts	16-19	Packets in the flow
dOctets	20-23	Octets (bytes) in the flow
first	24-27	SysUptime at start of the flow
last	28-31	SysUptime at the time the last packet of the flow was received

srcport	32-33	Layer 4 source port number or equivalent
dstport	34-35	Layer 4 destination port number or equivalent
pad1	36	Unused (zero) byte
tcp_flags	37	Cumulative OR of TCP flags
prot	38	Layer 4 protocol (for example, 6=TCP, 17=UDP)
tos	39	IP type-of-service byte
src_as	40-41	Autonomous system number of the source, either origin or peer
dst_as	42-43	Autonomous system number of the destination, either origin or peer
src_mask	44	Source address prefix mask bits
dst_mask	45	Destination address prefix mask bits
pad2	46-47	Pad 2 is unused (zero) bytes

Figure 5 is a typical flow record for the version 9 export format. As you can see NetFlow version 9 is different from the traditional NetFlow fixed format export record. In NetFlow version 9 a template describes the NetFlow data and the flow set contains the actual data. This allows for flexible export. Detailed information about the fields currently available in version 9 and export format architecture are available in the NetFlow Version 9 Flow-Record Format document.

Figure 5 NetFlow Version 9 Export Packet Example

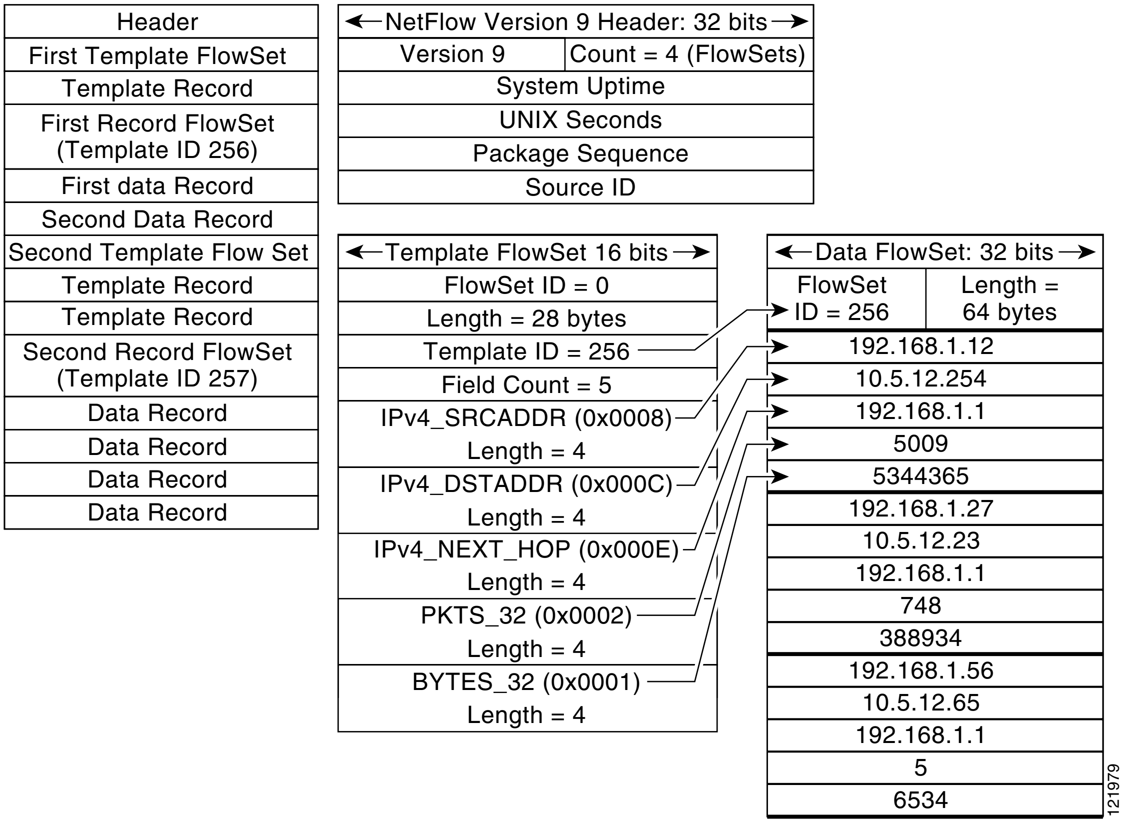

NetFlow data export packets are sent to a user-specified destination, such as the workstation running FlowCollector, either when the number of recently expired flows reaches a predetermined maximum, or every second—whichever occurs first. For a Version 1 datagram, up to 24 flows can be sent in a single UDP datagram of approximately 1200 bytes; for a Version 5 datagram, up to 30 flows can be sent in a single UDP datagram of approximately 1500 bytes; and for a Version 7 datagram, up to 28 flows can be sent in a single UDP datagram of approximately 1500 bytes.

Detailed information on the flow record formats, data types, and export data fields for version 1, 7, and 9 and platform specific information when applicable is shown in Appendix 2: Details for NetFlow Export Formats.

Aging NetFlow Cache Entries on a Routing Device

The routing device checks the NetFlow cache once per second and expires the flow in the following instances:

- Transport is completed (TCP FIN or RST).
- The flow cache has become full.
- The inactive timer has expired after 15 seconds of traffic inactivity.
- The active timer has expired after 30 minutes of traffic activity.

Setting NetFlow Active and Inactive Timers on a Routing Device

On a Cisco routing device, the following are default values of active and inactive timers:

- The inactive timer exports a packet with a default setting of 15 seconds of traffic inactivity. You can configure your own time interval for the inactive timer between 10 and 600 seconds.
- The active timer exports a packet after a default setting of 30 minutes of traffic activity. You can configure your own time interval for the active timer between 1 and 60 minutes.

Figure 6 illustrates how flow AT1 expires because the active timer for the flow exceeds the default value of 30 minutes. AT2 is the second flow which expires because the inactive timer exceeds the default value of 15 seconds.

Figure 6 Active and Inactive timers

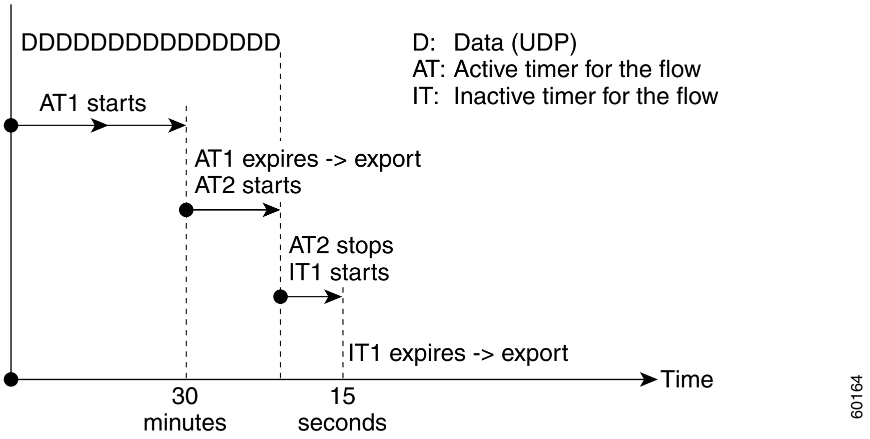

Catalyst 65k/7600 Flow Aging Timers

Catalyst switches use flow aging timers configured in a Multi-layer switching (MLS) cache, not the NetFlow cache used on routing devices. On a Catalyst switch, the following are default values of flow aging:

- Aging time: 256 seconds
- Fast aging time: disabled
- Long aging timer: 1,920 seconds

When flows expire from the NetFlow cache on a Catalyst switch, the flows are not exported. Catalyst switches export only when the export packet is full with 27 flows. Its default is 1920 seconds. Normal aging of a flow occurs when no more packets are switched for that flow for a predefined amount of time.

The normal aging table entries are purged when flow information has not been active for a user-configurable age time. Some users tune default timers if the cache is becoming full. Fast Aging can be used to reduce the Netflow table for short-duration connections that are already torn down and that therefore represent old information in the cache. For example, as an initial setting, the fast aging time might be configured to 128 seconds. That would ensure that short-lived flows or very slow flows would get aggressively purged. This type of change can help in reducing the growth of the Netflow table utilization when the number of flows is still well below the recommended upper bound and its trend of growth is low. A much more aggressive aging must instead be used when the Netflow table utilization gets closer and closer to its limit. You can use minimum fast aging time as the most aggressive way of purging active entries to make space for new flows. However, this drastic but sometimes necessary approach has the downside of increasing CPU utilization.

Cisco IOS Router-Based NetFlow Aggregation

Customers can expect a large volume of export data from NetFlow when it is enabled on many interfaces on high-end routers that switch many flows per unit time (such as the Cisco 12000 and Cisco 7500 Series). Designed to significantly reduce NetFlow Export data volume and improve NetFlow scalability, router-based NetFlow aggregation is a Cisco IOS software feature enhancement that enables router-based aggregation of NetFlow Export data. The eleven router-based NetFlow aggregation schemes enable the user to summarize NetFlow export data on the router before the data is exported to a NetFlow data collection device. With this feature enabled, aggregated NetFlow Export data is exported to a Collection device, resulting in lower bandwidth requirements for NetFlow Export data and reduced platform requirements for NetFlow data collection devices. Router based aggregation can be used with NetFlow Export Version 8 (v8) and Version 9 (v9).

The Router-based NetFlow Aggregation feature enables on-board aggregation by maintaining one or more extra NetFlow caches with different combinations of fields that determine which traditional flows are grouped together. These extra caches are called aggregation caches. As flows expire from the main flow cache, they are added to each enabled aggregation cache. The normal flow ager process runs on each active aggregation cache the just as it runs on the main cache. On demand aging is also supported. Figure 7 shows and example of how the main NetFlow cache can be aggregated into multiple aggregation caches based upon user configured aggregation schemes.

Figure 7 Building a NetFlow Aggregation Cache

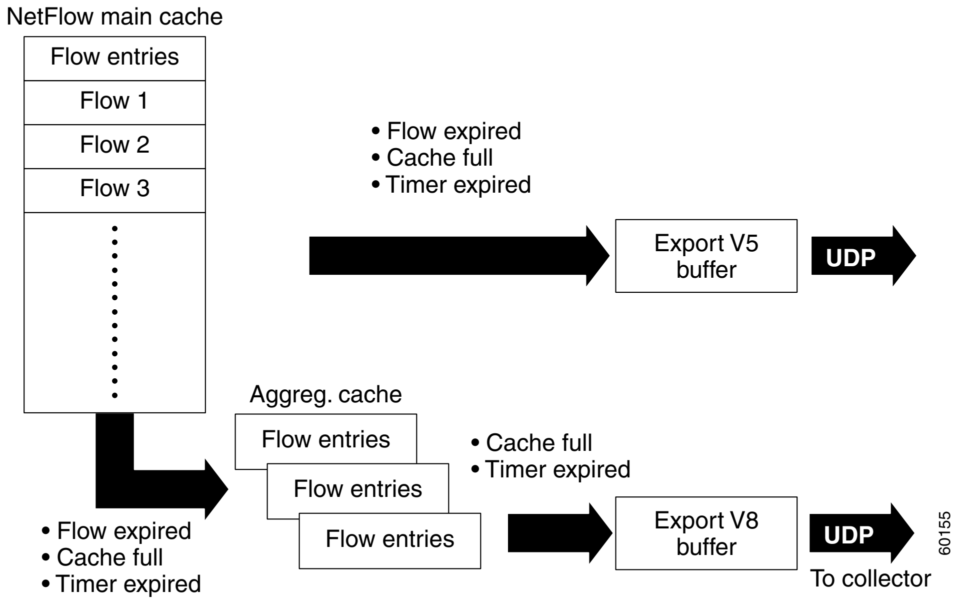

Cisco IOS Router-Based Aggregation with NetFlow v8 is available on all Cisco router platforms that support NetFlow beginning in releases 12.0(3)T and 12.0(3)S. NetFlow version 9 is available in IOS releases 12.3(1), 12.0(24)S, 12.2(18)S.

The default size for each secondary NetFlow aggregation cache (exported with v8 NetFlow Export datagrams) is 4096 entries on all platforms that support Cisco IOS NetFlow.

Use of Router-Based NetFlow Aggregation does not preclude the use of traditional NetFlow Services utilizing NetFlow Export v5 or v9. Router-Based NetFlow Aggregation (utilizing v8/v9 NetFlow Export datagrams) and traditional NetFlow Services (utilizing v9/v5 NetFlow Export datagrams) may be enabled simultaneously.

Selecting a NetFlow Aggregation Cache Scheme

You can configure each aggregation cache scheme with its individual cache size, cache ager timeout parameter, export destination IP address, and export destination UDP port. As data flows expire in the main cache (depending on the aggregation scheme configured), relevant information is extracted from the expired flow and the corresponding flow entry in the aggregation cache is updated. Each aggregation cache contains different field combinations that determine which data flows are grouped. The default aggregation cache size is 4096. The following are the 5 non-TOS based aggregation schemes:

- AS Aggregation Scheme
- Destination-Prefix Aggregation Scheme
- Prefix Aggregation Scheme

- Protocol-Port Aggregation Scheme
- Source Prefix Aggregation Scheme

The NetFlow ToS-Based Router Aggregation feature introduces support for six aggregation cache schemes that include the ToS byte as a field. The NetFlow ToS-Based Router Aggregation feature provides the ability to enable limited router-based ToS aggregation of NetFlow data, which results in summarized NetFlow data to be exported to a collection device. The following are the 6 TOS based aggregation schemes:

- AS-ToS Aggregation Scheme
- Destination-Prefix-ToS Aggregation Scheme
- Prefix-ToS Aggregation Scheme
- Protocol-Port-ToS Aggregation Scheme
- Source Prefix-ToS Aggregation Scheme
- Prefix-Port Aggregation Scheme

Tables 4 and 5 outline the router based aggregation Flow Record contents information.

Table 4 shows the Flow fields used in the non-TOS based aggregation schemes.

Table 4 Fields Used in the Non-TOS Based Aggregation Schemes

Field	AS	Protocol Port	Source Prefix	Destination Prefix	Prefix
Source Prefix			X		X
Source Prefix Mask			X		X
Destination Prefix				X	X
Destination Prefix Mask				X	X
Source App Port		X			
Destination App Port		X			
Input Interface	X		X		X
Output Interface	X			X	X
IP Protocol		X			
Source AS	X		X		X
Destination AS	X			X	X
First Timestamp	X	X	X	X	X
Last Timestamp	X	X	X	X	X
Number of Flows	X	X	X	X	X
Number of Packets	X	X	X	X	X
Number of Bytes	X	X	X	X	X

Table 5 shows the Flow fields used in the TOS based aggregation schemes.

Table 5 Flow fields used in the TOS based aggregation schemes

Field	AS-TOS	Protocol Port-TOS	Source Prefix-TOS	Destination Prefix-TOS	Prefix-TOS	Prefix-Port
Source Prefix			X		X	X
Source Prefix Mask			X		X	X
Destination Prefix				X	X	X
Destination Prefix Mask				X	X	X
Source App Port		X				X
Destination App Port		X				X
Input Interface	X	X	X		X	X
Output Interface	X	X		X	X	X
IP Protocol		X				X
Source AS	X		X		X	
Destination AS	X			X	X	
TOS	X	X	X	X	X	X
First Timestamp	X	X	X	X		X
Last Timestamp	X	X	X	X		X
Number of Flows	X	X	X	X		X
Number of Packets	X	X	X	X		X
Number of Bytes	X	X	X	X		X

Table 6 lists the number of flows in a UDP datagram packet and the packet length (in bytes) for the various export version formats.

Table 6 Flows and Packet Lengths for all NetFlow Export Versions

NetFlow Export Version Format	Number of Flows in a Packet	Packet Length (bytes)
V1	24	Approx. 1,200
V5	30	Approx. 1,500
V7	27	Approx. 1,500
V8 AsMatrix	51	1456
V8 Protocol-PortMatrix	51	1456
V8 Source-PrefixMatrix	44	1436
V8 Destination-PrefixMatrix	44	1436
V8 PrefixMatrix	35	1428
V8 As-ToSMatrix	44	1436
V8 Protocol-Port-ToSMatrix	44	1436
V8 Source-Prefix-ToSMatrix	44	1436
V8 Destination-PrefixMatrix	44	1436

NetFlow Export Version Format	Number of Flows in a Packet	Packet Length (bytes)
V8 Prefix-ToSMatrix	35	1428
V8 Prefix-PortMatrix	35	1428

For more information on router based aggregation see Appendix 3: Router Based Aggregation Schemes And Detailed NetFlow Export Formats for detailed export formats details. Also the IOS documentation NetFlow Aggregation and NetFlow ToS-Based Router Aggregation Feature Overview has more information

Catalyst 65k Flow Mask information

Flow keys are a set of values that determine how a flow is identified. Typically on most Cisco devices the flow keys are a fixed 7 tuple of information. The Catalyst 65k has the capability to define a Flow Mask, which is a predefined set of flow key values that is configured by the user. The Flow Masks will perform automatic aggregation of data in the NetFlow cache. So for example, if the user is interested in accounting for packets from the same source IP address going to the same destination IP address and aggregating this traffic into one flow then they can use the destination-source Flow mask (see Figure 8 below). This concept of flow mask is different than an aggregation scheme in which aggregation of the data takes place after the complete set of 7 flow keys is used to create the flow information. With a flow mask the flow information is aggregated directly into the main MLS (NetFlow) cache on the 65k. The main reason to use the Flow Mask feature is to enhance scalability by utilizing the NetFlow cache efficiently, aggregate flows and decrease the amount of flow export. While a flow mask does increase efficient use of the NetFlow cache the amount of detailed information is reduced with the aggregation of the data flows. Figure 8 shows the Catalyst 65k/7600 Flow Masks.

Figure 8 Catalyst 65k/7600 Flow Masks

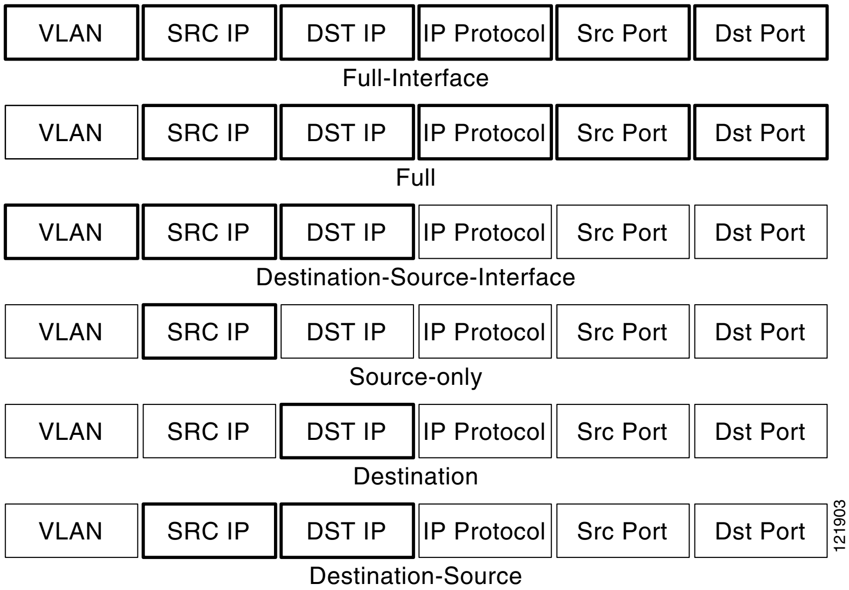

Export Version Information for Cisco Platforms

Table 7 outlines the first releases for specific NetFlow export versions per platform. For specific feature information and release and platform support use Feature Navigator

Table 7 First Releases for NetFlow Export Versions by Platform

Cisco IOS Software Release Version	NetFlow Export Version(s)	Supported Cisco Hardware Platforms
11.1CA	v1, v5	Cisco 7200, 7500 were the first platforms in 11.1CA. v5 is now available for all IOS platforms.
12.3(1),12.0(24)S,12.2(18)S, 12.3(2)T	v9	Cisco 800, 1700, 2600, 3600,3700,6400,7200,7300,7500,12000
12.0(14)S	v5	Cisco 12000
12.0(6)S	v8	Cisco 12000
See Table 6	v5,v7,v8	Catalyst 65k
12.1(13)EW	v5	Catalyst 4k Supervisor 4
12.1(19)EW	v8	Catalyst 4k Supervisor 4
12.1(18)EW	v5,v8	Catalyst 4k Supervisor 5

Table 8 shows the Catalyst 65k/7600 NetFlow version Support

Table 8 Catalyst 65k/7600 NetFlow version Support

Supervisor	Hybrid	Native 12.1E	Native 12.2SX
MSFCx	v5	v5	v5, v8*
Sup1a	V7, v8	v7	N/A
Sup2	V7, v8	v5, v7	v5, v7, v8
Sup720	v5, v7, v8	v5, v7	v5, v7, v8

NetFlow Performance Information

A specific white paper has been written to give details of how NetFlow implicates performance on software based Cisco platforms. NetFlow performance impact comes mainly from the characterization of the flow information in the NetFlow cache and the formation of the NetFlow export packet and the export process. In general NetFlow is supported in hardware ASIC on many Cisco platforms including the Catalyst 4500, 6500, 7600, 10000 and 12000 routers. When NetFlow is utilized in hardware the main performance impact is due to export of the flow information but the characterization of the flows is done in hardware.

The export version does not affect performance numbers for NetFlow including v5, v8 or v9.

The additional CPU utilization on software platforms due to NetFlow varies based on the number of flows.

Table 9 shows the approximate CPU utilization for a number of active flows.

Table 9 Approximate CPU utilization for a number of active flows

Number of Active Flow in Cache	Additional CPU Utilization
10000	< 4%
45000	<12%
65000	<16%

Sampled Netflow will significantly decrease CPU utilization to the router. On average sampled NetFlow 1:1000 packets will reduce CPU by 82% and 1:100 sampling packets reduce CPU by 75% on software platforms. The conclusion is sampled NetFlow is a significant factor in reducing CPU utilization. See the section below in this document on sampled NetFlow for more information on sampling techniques used by Cisco devices.

In general dual export has no significant CPU impact on the router and this feature available in IOS 12.0(19)S, 12.2(2)T, 12.2(14)S for redundancy of the export.

Some significant factors in reducing CPU utilization from the NetFlow process include:

- Sampled NetFlow
- Optimize the aging timers to proper values for the amount of flows
- Leverage a distributed architecture
- Utilization of flow masks on Catalyst 65k/7600

Please see the NetFlow Performance Analysis white paper for more information.

NetFlow Memory Allocation Information:

The NetFlow cache size can vary from 1k to 512K and is configurable for software based platforms such as 75xx and 72xx. Each Cache entry consumes about 64 bytes of memory. The amount of memory on a Cisco 12K line card denotes how many flows are possible in the cache. For example, if an engine 3 line card has 256M bytes of memory, NetFlow allocates 256M/16/64=256k entries. If NetFlow aggregation (discussed later) is used then depending on user configuration, up to 512K entries are possible. The Cisco Catalyst 65k/7600 will have different effective hardware cache sizes based on the supervisor card and PFC.

Table 10 shows the Catalyst 65k Hardware Cache Effective Sizes.

Table 10 Catalyst 65k Hardware Cache Effective Sizes

Catalyst 65k/7600 PFC	Effective Number of NetFlow Cache Entries Available
PFC2/DFC	32K entries
PFC3A/DFC3A	64K entries
PFC3B/DFC3B	115K entries
PFC3BXL/DFC3BXL	230K entries

The number of cache entry changes per PRE on the Cisco 10000 router is shown in Table 11.

Table 11 Cisco 10000 NetFlow Cache Sizes

Cisco 10000	NetFlow Entries Available
PRE1	512K entries
PRE2	1M entries

Sampled NetFlow Details and Platform Support

Cisco was the first company to implement packet sampling for NetFlow on the Cisco 12000 router. On an interface, Sampled NetFlow allows you to collect NetFlow statistics for a subset of incoming (ingress) traffic. Sampled NetFlow significantly reduces CPU utilization on a router, reduces export volume but still allows a view of most IP flows switching in the device. Sampling is very useful for capacity planning or network planning when every flow may not be needed to understand the network behavior. There are 3 types of sampling used on Cisco platforms: deterministic sampling, time based sampling and random sampling. Deterministic sampling will select every Nth packets, with N specified by the user. Random sampling will randomly select one out of N packets with N specified by the user. Time based sampling will select a sampled packet every N milli-seconds. Random sampling is considered the best technique for packet sampling.

Figure 9 contrasts deterministic and random sampling.

Figure 9 Deterministic and Random Sampling

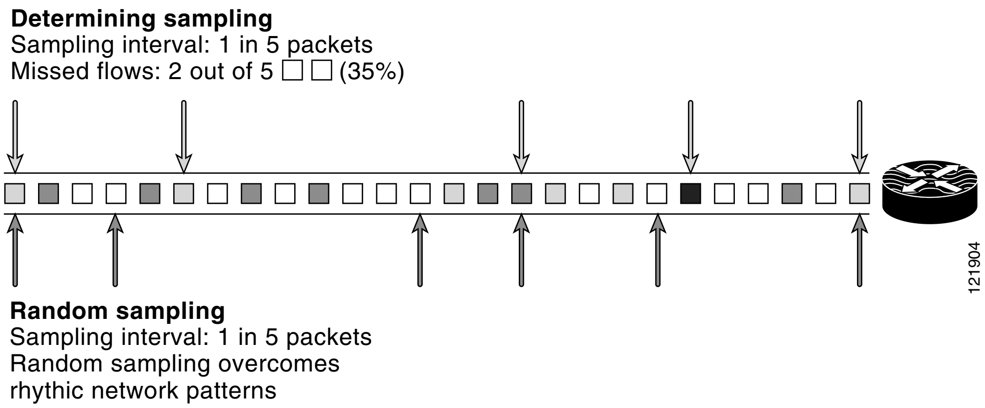
Determining sampling
Sampling interval: 1 in 5 packets
Missed flows: 2 out of 5 □ □ (35%)
Random sampling
Sampling interval: 1 in 5 packets
Random sampling overcomes
rhythic network patterns
121904

Table 12 describes when each type of sampling was first introduced per platform and the version of NetFlow export supported.

Table 12 Sampling by Type, Version of NetFlow Export Supported, and IOS Version

Platform Name	Sampling Type	NetFlow Export Version(s)	IOS Release
Software based platforms	Random Sampling	All	12.0(26)S, 12.2(18)S, 12.3(2)T
Cisco 12000	Deterministic Sampling	v5,v8,v9	See Table 13 Below
Catalyst 65k/7600	Random Sampling	v5,v8,v7	12.1(13)E
Catalyst 65k/7600	Time based Sampling	v5,v8,v7	12.1(13)E

On the Cisco 12000 platform the introduction of sampled NetFlow varies per line card.

Table 13 shows line card support for sampled and Full (non-sampled) NetFlow.

Table 13 Line card support for sampled and Full (non-sampled) NetFlow

Engine	"Full" NetFlow	Sampled NetFlow
0	Supported	Supported
1	Supported	Supported
2		Supported
3	Aggregated Only v8	Supported
4		
4+		Supported

Table 14 Shows Cisco 12000 sampled NetFlow and Full NetFlow release information per IOS release.

Table 14 Cisco 12000 sampled NetFlow and Full NetFlow release information per IOS Release

Field	Full NetFlow		Sampled NetFlow	
	Version 5	Version 8	Version 5	Version 8

Engine 0	12.0(14)S	12.0(6)S	12.0(14)S	12.0(11)S
Engine 2 POS	N/A	N/A	12.0(14)S	12.0(14)S
Engine 2 3xGE	N/A	N/A	12.0(16)S	12.0(16)S
Engine 3	N/A	12.0(21)S	12.0(21)S	12.0(21)S
Engine 4	N/A	N/A	N/A	N/A
Engine 4+ POS	N/A	N/A	12.0(21)S	12.0(21)S
Ashara	N/A	N/A	12.0(21)S	12.0(21)S
Tango	N/A	N/A	12.0(21)S	12.0(21)S

For more information on sample NetFlow see the IOS documentation on Random Sampled NetFlow.

Supported Interfaces, Encapsulations and Protocols

NetFlow supports IPv4 (and IPv4-encapsulated) routed traffic over a wide range of interface types and encapsulations. This includes Frame Relay, Asynchronous Transfer Mode, Inter-Switch Link, 802.1q, Multi-link Point to Point Protocol, General Routing Encapsulation, Layer 2 Tunneling Protocol, Multi-protocol Label Switching VPNs, and IP Sec Tunnels. For detailed information on encapsulation types supported and tested see the NetFlow on Logical Interfaces white paper.

To account for traffic entering a tunnel specify generic ingress NetFlow on the router. To account for tunnel and post tunnel flows NetFlow can be configured on the tunnel interface at the tunnel end point. A white paper has been written about NetFlow for GRE and IPSec tunnels.

NetFlow is supported per sub-interface. If NetFlow is configured on the major interface then all sub-interfaces will be accounted. Also available is NetFlow Subinterface feature to account for packets on specific sub-interfaces.

NetFlow support for multicast does exist on some Cisco platforms. For more information on Multicast NetFlow see the Cisco IOS Multicast NetFlow documentation.

NetFlow supports IPv6 environments in 12.3(7)T and above. For more information on IPv6 and netFlow see the Cisco IOS IPv6 NetFlow documentation.

NetFlow can be used effectively in an MPLS network for VPN accounting or capacity planning. Generic ingress NetFlow can be used to account for traffic from the customer site entering an MPLS network per VPN. The customer name can be correlated to the VRF associated with the particular customer site. Two other features specifically designed for an MPLS network include MPLS egress NetFlow and MPLS Aware NetFlow. These features are available on 3700, 3800, 7200, 7500 and Cisco 12000 series routers. Egress NetFlow Accounting will account for packets leaving an MPLS cloud and egress to a specific customer site. This feature is useful for VPN accounting. MPLS aware NetFlow is used on MPLS core routers to account for traffic and aggregate traffic per MPLS label. This feature effectively tells the user how much traffic is destined for a specific PE router in the network, allowing the user to calculate a traffic matrix between PE routers for the MPLS network.

NetFlow and Quality of Service (QoS)

NetFlow records contain the Type of Service (ToS) field in the IP header as well as application ports, traffic volumes and timestamps. This allows the user to understand traffic profiles per class of service (COS) for WAN traffic including data, voice and video. The user of NetFlow can verify the QOS levels achieved and optimize bandwidth for specific classes of service.

NetFlow Activation and Deployment Information

Cisco recommends careful planning of NetFlow deployment with NetFlow services activated on strategically located edge/aggregation routers which capture the data required for planning, monitoring and accounting applications. Key deployment considerations include the following:

- Understanding your application-driven data collection requirements: accounting applications may only require originating and terminating router flow information whereas monitoring applications may require a more comprehensive (data intensive) end-to-end view
- Understanding the impact of network topology and routing policy on flow collection strategy: for example, avoid collecting duplicate flows by activating NetFlow on key aggregation routers where traffic originates or terminates and not on backbone routers or intermediate routers which would provide duplicate views of the same flow information
- NetFlow can be implemented in CLI to understand the number of flow in the network and the impact on the router. NetFlow export can then be setup at a later date to complete the NetFlow deployment.

NetFlow is in general an ingress measurement technology which should be deployed on appropriate interfaces on edge/aggregation or WAN access routers to gain a comprehensive view of originating and terminating traffic to meet customer needs for accounting, monitoring or network planning data. Egress NetFlow accounting is available in the latest release of IOS including 12.3(11)T and also output NetFlow on the GSR engine 3 line cards. The key mechanism for enhancing NetFlow data volume manageability is careful planning of NetFlow deployment. NetFlow can be deployed incrementally (i.e. interface by interface) and strategically (i.e. on well chosen routers) —instead of widespread deployment of NetFlow on every router in the network. Cisco will work with customers to determine key routers and key interfaces where NetFlow should be activated based on the customer's traffic flow patterns and network topology and architecture.

NetFlow Management Applications

Cisco IOS NetFlow software is part of a larger family of products, management utilities and partner applications designed to gather and export flow statistics, collect and perform data volume reduction on the exported statistics, and feed flow detail records to consumer applications such as planning, accounting and monitoring. Figure 10 is an example of a Typical NetFlow Infrastructure.

Figure 10 NetFlow Infrastructure

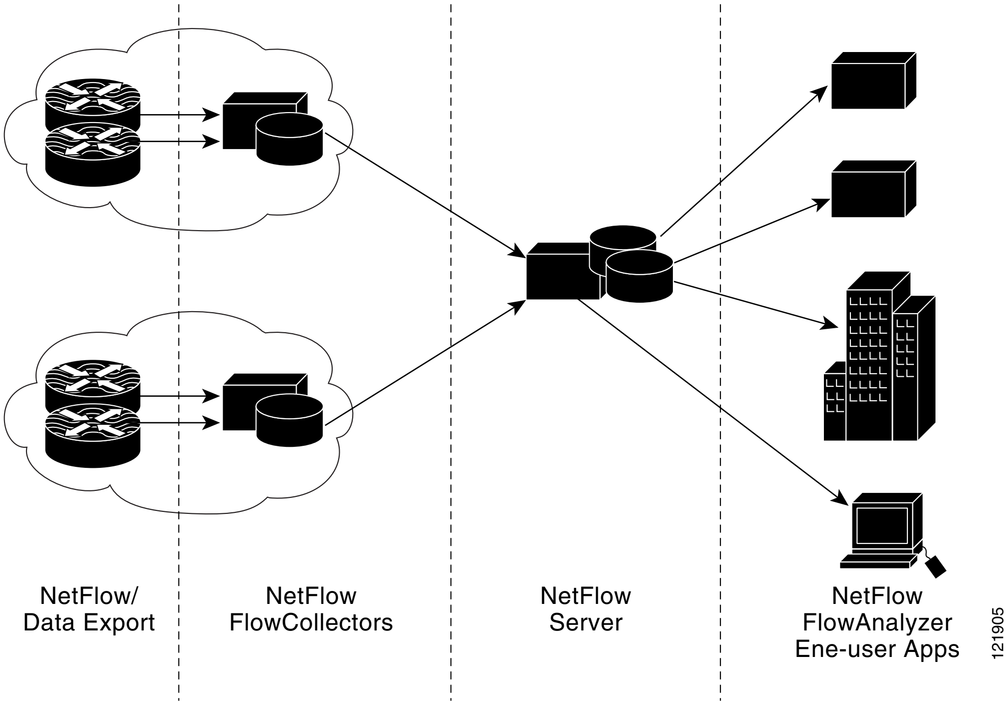

NetFlow management applications:

- Collect, store and perform data volume reduction on exported NetFlow data
- Provide a scalable and distributed NetFlow data collection and consolidation architecture
- Provide network monitoring, analysis, and troubleshooting tools

NetFlow Collectors

Cisco NetFlow collector provides fast, scalable, and economical data collection from multiple NetFlow Export-enabled devices. The Cisco collector consumes flow datagrams from multiple NetFlow Export-enabled devices and performs data volume reduction through selective filtering and aggregation, performs bi-directional flow analysis and flow de-duplication. The Cisco Network Analysis Module (NAM) can collect NetFlow data within Cisco devices and provides a comprehensive reporting and traffic analysis solution. Other third party traffic analysis, billing, security, and monitoring applications are available for NetFlow. For a complete list of Cisco partners go to the NetFlow partner web page.

NetFlow MIB

The NetFlow MIB is a new interface for NetFlow data now available on Cisco routers in 12.3(7)T and 12.2(25)S. The NetFlow MIB can be used to access NetFlow data when export is not possible from a specific network location or for troubleshooting network behavior. For instance the NetFlow MIB can be used to detect security violations within the network. The NetFlow MIB is not a replacement for traditional export. This feature allows the user to configure NetFlow commands using SNMP and retrieve essential flow information including number of flows, flow per second per protocol, packets and bytes per flow. Also included with the NetFlow MIB and CLI is the Top Talkers feature to show users the Top N Flows available in the NetFlow cache. For more information on the NetFlow MIB and MIB with Top Talkers see the following links:

NetFlow MIB

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_feature_guide09186a008020df09.html

NetFlow MIB and Top Talkers

http://www.cisco.com/en/US/products/sw/iosswrel/ps1838/products_feature_guide09186a0080259533.html

NetFlow IOS Configuration

Enabling NetFlow on a Router Interface

Using the **ip route-cache flow** command, you can enable NetFlow on an interface. If you enable NetFlow on an interface that contains subinterfaces, all the subinterfaces will be enabled automatically.

The remainder of this section contains sixteen configuration examples. There are eleven example configurations for Cisco routers. These are followed by five example configurations for Catalyst 65k switches and 7600 Routers.

Configuration Example 1: Configure NetFlow Export Destination Using Version 9

Configuration Example 2: Configuring an AS Aggregation Cache Scheme

Configuration Example 3: Configuring an AS-ToS Aggregation Cache Scheme

Configuration Example 4: Setting an Active and Inactive Timer for a NetFlow Cache

Configuration Example 5: Setting an Active and Inactive Timer for a NetFlow Aggregation Cache

Configuration Example 6: Configuring NetFlow Sampled Mode and the Packet Interval

Configuration Example 7: NetFlow Multiple Export Destinations

Configuration Example 8: NetFlow Minimum Source and Destination Mask for Prefix Aggregation Cache

Configuration Example 9: NetFlow Minimum Destination Mask for Destination-Prefix Aggregation Cache—Configuration Example

Configuration Example 10: NetFlow Multiple Export Destinations on an Aggregation Cache

Configuration Example 11: Exporting a Peer AS Packet Using Version 5 Format

Configuration Example 1: Configure NetFlow Export Destination Using Version 9

With the Version 9 format, you typically export the Version 9 record format when the NetFlow cache is full, the flow has expired, or the timer has expired. The following configuration example shows how to configure an export destination using the Version 9 record format:

In this example, you configure the routing device to export Version 5 datagram format NetFlow cache entries to a workstation using the using the **ip flow-export version 9** command.

The **ip flow-export source loopback interface** command causes the router to use the IP address assigned to the loopback interface as the source IP address for the NetFlow packets (the NFC will collect exported packets from this routing device's source IP address). If you do not configure the **ip flow-export source loopback** command and the exported packet takes another outgoing interface, the source IP address will change and the collection device will interpret this exported packet as coming from a different routing device.

```
Router(config-if)# ip route-cache flow
Router(config)# ip flow-export destination 172.17.246.225 9996
Router(config)# ip flow-export version 9
Router(config)# ip flow-export source loopback 0
```

Verifying the Statistics of a NetFlow Cache Using the Version 9 Format

In the output of the **show ip cache flow** command below, a summary of the NetFlow cache statistics using the **show ip cache flow** command is displayed in EXEC mode. For more accounting statistics, you can use the **verbose** keyword, using the **show ip cache verbose flow** command in EXEC mode.

```
Router# show ip cache flow

IP packet size distribution (12718M total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .554 .042 .017 .015 .009 .009 .009 .013 .030 .006 .007 .005 .004 .004

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .003 .007 .139 .019 .098 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 4456448 bytes
  65509 active, 27 inactive, 820628747 added
  955454490 ager polls, 0 flow alloc failures
  Exporting flows to 1.1.15.1 (2057)
  820563238 flows exported in 34485239 udp datagrams, 0 failed
  last clearing of statistics 00:00:03

Protocol         Total  Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
                 Flows   /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-Telnet     2656855    4.3        86    78    372.3      49.6      27.6
TCP-FTP        5900082    9.5         9    71     86.8      11.4      33.1
TCP-FTPD       3200453    5.1       193   461   1006.3      45.8      33.4
TCP-WWW      546778274  887.3        12   325  11170.8       8.0      32.3
TCP-SMTP      25536863   41.4        21   283    876.5      10.9      31.3
TCP-X           116391    0.1       231   269     43.8      68.2      27.3
TCP-BGP          24520    0.0        28   216      1.1      26.2      39.0
TCP-Frag         56847    0.0        24   952      2.2      13.1      33.2
TCP-other     49148540   79.7        47   338   3752.6      30.7      32.2
UDP-DNS      117240379  190.2         3   112    570.8       7.5      34.7
UDP-NTP        9378269   15.2         1    76     16.2       2.2      38.7
UDP-TFTP          8077    0.0         3    62      0.0       9.7      33.2
UDP-Frag         51161    0.0        14   322      1.2      11.0      39.4
UDP-other     45502422   73.8        30   174   2272.7       8.5      37.8
ICMP          14837957   24.0         5   224    125.8      12.1      34.3
IGMP             40916    0.0       170   207     11.3     197.3      13.5
IPINIP            3988    0.0     48713   393    315.2     644.2      19.6
GRE               3838    0.0        79   101      0.4      47.3      25.9
```

```
IP-other         77406    0.1        47   259       5.9       52.4      27.0
Total:       820563238 1331.7        15   304  20633.0        9.8      33.0

SrcIf    SrcIPaddress    DstIf    DstIPaddress    Pr SrcP DstP Pkts B/Pk Active
Fd0/0    80.0.0.3        Hs1/0    200.1.9.1       06 0621 0052    7   87    5.9
Fd0/0    80.0.0.3        Hs1/0    200.1.8.1       06 0620 0052    7   87    1.8
Hs1/0    200.0.0.3       Fd0/0    80.1.10.1       06 0052 0621    6   58    1.8
Hs1/0    200.0.0.3       Fd0/0    80.1.1.1        06 0052 0620    5   62    5.9
Fd0/0    80.0.0.3        Hs1/0    200.1.3.1       06 0723 0052   16   68    0.3
HS1/0    200.0.0.3       Fd0/0    80.1.2.1        06 0052 0726    6   58   11.8
Fd0/0    80.0.0.3        Hs1/0    200.1.5.1       06 0726 0052    6   96    0.3
Hs1/0    200.0.0.3       Fd0/0    80.1.4.1        06 0052 0442    3   76    0.3
Hs1/0    200.0.0.3       Fd0/0    80.1.7.1        06 0052 D381   11 1171    0.6
```

In this show output, a summary of the NetFlow cache statistics using the **show ip cache flow** command is displayed in EXEC mode. For more accounting statistics, you can use the **verbose** keyword, using the **show ip cache verbose flow** command in EXEC mode

The following **show ip cache verbose flow** output displays NetFlow cache verbose statistics including the entire content of the NetFlow cache, which will be exported to the NFC after the flow expires:

```
Router# show ip cache verbose flow

IP packet size distribution (12718M total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .554 .042 .017 .015 .009 .009 .009 .013 .030 .006 .007 .005 .004 .004

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .003 .007 .139 .019 .098 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 4456448 bytes
  65509 active, 27 inactive, 820628747 added
  955454490 ager polls, 0 flow alloc failures
  Exporting flows to 1.1.15.1 (2057)
  820563238 flows exported in 34485239 udp datagrams, 0 failed
  last clearing of statistics 00:00:03

Protocol         Total  Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
                 Flows   /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-Telnet     2656855    4.3        86    78    372.3      49.6      27.6
TCP-FTP        5900082    9.5         9    71     86.8      11.4      33.1
TCP-FTPD       3200453    5.1       193   461   1006.3      45.8      33.4
TCP-WWW      546778274  887.3        12   325  11170.8       8.0      32.3
TCP-SMTP      25536863   41.4        21   283    876.5      10.9      31.3
TCP-X           116391    0.1       231   269     43.8      68.2      27.3
TCP-BGP          24520    0.0        28   216      1.1      26.2      39.0
TCP-Frag         56847    0.0        24   952      2.2      13.1      33.2
TCP-other     49148540   79.7        47   338   3752.6      30.7      32.2
UDP-DNS      117240379  190.2         3   112    570.8       7.5      34.7
UDP-NTP        9378269   15.2         1    76     16.2       2.2      38.7
UDP-TFTP          8077    0.0         3    62      0.0       9.7      33.2
UDP-Frag         51161    0.0        14   322      1.2      11.0      39.4
UDP-other     45502422   73.8        30   174   2272.7       8.5      37.8
ICMP          14837957   24.0         5   224    125.8      12.1      34.3
IGMP             40916    0.0       170   207     11.3     197.3      13.5
IPINIP            3988    0.0     48713   393    315.2     644.2      19.6
GRE               3838    0.0        79   101      0.4      47.3      25.9
IP-other         77406    0.1        47   259      5.9      52.4      27.0
Total:       820563238 1331.7        15   304  20633.0       9.8      33.0

SrcIF        SrcIPaddress    DstIf        DstIPaddress    Pr TOS Flgs Pkts
port Msk AS                  Port Msk AS  NextHop                B/Pk Active
Se0/1        193.1.1.3       Se0/0        172.17.246.228  11 00  10   5
00A1 /24 193                 C628 /0 0    0.0.0.0                84   39.7
```

In this show output, a summary of the NetFlow cache statistics using the **show ip cache verbose** flow command is displayed in EXEC mode.

Configuration Example 2: Configuring an AS Aggregation Cache Scheme

By maintaining one or more extra flow caches, called *aggregation caches*, the Router-Based Aggregation feature allows limited aggregation of NDE streams on a routing device.

The following configuration example shows how to configure an export destination using the Version 8 or 9 record format:

```
Router(config)# ip flow-aggregation cache as
Router(config-flow-cache)# export destination 172.17.246.225 9996
Router(config-flow-cache)# enabled
```

In this example, you enable an AS aggregation cache using the **ip route aggregation cache as** command. You enable data export of the Version 8 or Version 9 datagram format NetFlow cache entries to a workstation using the export destination command. Using the enabled command, you enable aggregation cache creation

Verifying Statistics of an AS Aggregation Cache

The output of the **show ip cache flow aggregation as** command below displays the statistics of an AS aggregation cache:

```
Router# show ip cache flow aggregation as

IP Flow Switching Cache, 278544 bytes
  2 active, 4094 inactive, 13 added
  178 ager polls, 0 flow alloc failures

Src If         Src AS  Dst If         Dst AS  Flows   Pkts  B/Pk  Active
Fa1/0           0      Null            0         1      2    49    10.2
Fa1/0           0      Se2/0          20         1      5   100     0.0
```

In this show output, a summary of the AS aggregation cache statistics using the **show ip cache flow aggregation** command is displayed in EXEC mode.

Configuration Example 3: Configuring an AS-ToS Aggregation Cache Scheme

The following configuration example shows how to configure an export destination using the Version 8 or 9 record format:

```
Router(config)# ip flow-aggregation cache as-tos
Router(config-flow-cache)# cache timeout inactive 120
Router(config-flow-cache)# export destination 2.2.2.2 3000
Router(config-flow-cache)# enabled
```

In this example, you enabled an AS-ToS aggregation cache using the ip flow-aggregation cache as-tos command. Using the cache timeout inactive 120 command, you configure the inactive timer to expire flows after 120 seconds of inactivity. You enable data export of the Version 8 or Version 9 datagram format NetFlow cache entries to a workstation using the export destination command. Using the enabled command, you enable aggregation cache creation.

Verifying Statistics of an AS-ToS Aggregation Cache

The output of the **show ip cache verbose flow aggregation as-tos** output below displays the statistics of an AS aggregation cache:

```
Router# show ip cache verbose flow aggregation as-tos

IP Flow Switching Cache, 278544 bytes
  4 active, 4094 inactive, 103 added
  1609 ager polls, 0 flow alloc failures

Src If       Src AS  Dst If       Dst AS  TOS  Flows   Pkts  B/Pk  Active
Et1/2         50     Fd4/0         40     CC    1      3568   28    17.8
Et1/2         0      Fd4/0         40     C0    15     17K    28    17.8
Et1/1         50     Fd4/0         40     55    1      3748   28    17.8
Et1/2         0      Null          0      C0    1      2      49     0.9
```

Configuring a NetFlow Cache Size

A NetFlow cache can be resized depending on the platform and the amount of DRAM on a line card. Using the **ip flow-cache entries** command, you can configure the size of your NetFlow cache between 1,024 entries and to a maximum of 524,288 entries. Using the **cache entries** command (after you configured NetFlow aggregation), you can configure the size of the NetFlow aggregation cache from 1,024 entries to a maximum of 524,288 entries. A NetFlow cache entry (a single flow) is 64 bytes

Note Cisco recommends you not change the default number of NetFlow cache entries. This could lead to improper use and cause network problems.

The changes to the NetFlow cache will take place after enabling or disabling NetFlow or after a reboot.

Configuration Example 4: Setting an Active and Inactive Timer for a NetFlow Cache

The following configuration example shows how to configure the amount of time till expiration of a flow using an active and inactive time for a NetFlow cache. This example does not apply to Catalyst 65k or 76xx which uses similar but different aging timers discussed earlier in this document.

```
Router(config)# ip flow-cache timeout active 20
Router(config)# ip flow-cache timeout inactive 130
```

In this example, you configure the active timer to expire flows after 20 minutes using the **ip flow-cache timeout active** command. Using the **ip flow-cache timeout inactive** command, you configure the inactive timer to expire flows after 130 seconds of inactivity.

The output of the the **show ip cache verbose flow** command below displays the configured expiration setting for active and inactive timers for a NetFlow cache:

```
Router# show ip cache verbose flow

IP packet size distribution (12718M total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .554 .042 .017 .015 .009 .009 .009 .013 .030 .006 .007 .005 .004 .004

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .003 .007 .139 .019 .098 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 4456448 bytes
  65509 active, 27 inactive, 820628747 added
  955454490 ager polls, 0 flow alloc failures
  Active flows timeout in 20 minutes
  Inactive flows timeout in 130 seconds
  Exporting flows to 1.1.15.1 (2057)
  820563238 flows exported in 34485239 udp datagrams, 0 failed
  last clearing of statistics 00:00:03

Protocol         Total  Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
                 Flows   /Sec     /Flow  /Pkt     /Sec     /Flow     /Flow
TCP-Telnet     2656855    4.3        86    78    372.3      49.6      27.6
TCP-FTP        5900082    9.5         9    71     86.8      11.4      33.1
TCP-FTPD       3200453    5.1       193   461   1006.3      45.8      33.4
TCP-WWW      546778274  887.3        12   325  11170.8       8.0      32.3
TCP-SMTP      25536863   41.4        21   283    876.5      10.9      31.3
TCP-X           116391    0.1       231   269     43.8      68.2      27.3
TCP-BGP          24520    0.0        28   216      1.1      26.2      39.0
TCP-Frag         56847    0.0        24   952      2.2      13.1      33.2
TCP-other     49148540   79.7        47   338   3752.6      30.7      32.2
UDP-DNS      117240379  190.2         3   112    570.8       7.5      34.7
UDP-NTP        9378269   15.2         1    76     16.2       2.2      38.7
UDP-TFTP          8077    0.0         3    62      0.0       9.7      33.2
UDP-Frag         51161    0.0        14   322      1.2      11.0      39.4
UDP-other     45502422   73.8        30   174   2272.7       8.5      37.8
ICMP          14837957   24.0         5   224    125.8      12.1      34.3
IGMP             40916    0.0       170   207     11.3     197.3      13.5
IPINIP            3988    0.0     48713   393    315.2     644.2      19.6
GRE               3838    0.0        79   101      0.4      47.3      25.9
IP-other         77406    0.1        47   259      5.9      52.4      27.0
Total:       820563238 1331.7        15   304  20633.0       9.8      33.0

SrcIF        SrcIPaddress    DstIf        DstIPaddress    Pr TOS Flgs Pkts
port Msk AS                  Port Msk AS  NextHop                B/Pk Active
Se0/1        193.1.1.3       Se0/0        172.17.246.228  11 00  10   5
00A1 /24 193                 C628 /0 0    0.0.0.0                84   39.7
```

In this show output, the active timer expires flows after 20 minutes and the inactive timer causes flows to expire after 130 seconds of inactivity. The "Flgs" output verifies that the TCP flag is 10.

Configuration Example 5: Setting an Active and Inactive Timer for a NetFlow Aggregation Cache

The following configuration example shows how to configure the amount of time till expiration of a flow using active and inactive timers for a NetFlow aggregation cache:

```
Router(config)# ip flow-aggregation cache prefix
Router(config)# cache timeout active 25
Router(config)# cache timeout inactive 400
```

In this example, you enable a prefix aggregation cache using the ip flow-aggregation cache prefix command. You configure the active timer to expire flows after 25 minutes using the cache timeout active command. Using the cache timeout inactive command, you configure the inactive timer to expire flows after 400 seconds of inactivity.

Verifying Active and Inactive Timers for a NetFlow Aggregation Cache

The output of the **show ip cache flow aggregate prefix** command below displays the configured expiration setting for active and inactive timers for a NetFlow aggregation cache:

```
Router# show ip cache flow aggregate prefix

ip Flow Switching Cache, 0 bytes
0active, 0 inactive, 0 added
0 ager polls, 0 flow alloc failures
Active flows timeout in 25 minutes
Inactive flows timeout in 400 seconds
```

In this show output, the active timer expires flows after 25 minutes and the inactive timer causes flow to expire after 400 seconds of inactivity.

Configuration Example 6: Configuring NetFlow Sampled Mode and the Packet Interval

The following configuration example shows how to configure NetFlow random sampled mode globally and apply the sampler map to the interface:

```
Router(config)# flow-sampler-map mysampler1
Router(config-sampler)# mode random one-out-of 100
Router(config-if)# interface FastEthernet9/0/0
Router(config-if)# mac-address 4000.0209.0000
Router(config-if)# ip address 150.1.2.2 255.255.255.0
Router(config-if)# no ip directed-broadcast
Router(config-if)# flow-sampler mysampler1
Router(config-if)# ip route-cache distributed
Router(config-if)# no ip mroute-cache
Router(config-if)# no keepalive
```

In this example, you configure NetFlow sampled mode on interface FastEthernet9/0/0 using the flow-sampler mysampler1 command.

The output of the **show flow-sampler** command displays Sampled NetFlow settings for the sampling mode and sampling interval:

```
Router#show flow-sampler

 Sampler : my, id : 1, packets matched : 0, mode : random sampling mode
  sampling interval is : 100
```

In this **show** output, you can verify that Random Sampled NetFlow is enabled and the packet interval sampled mode is configured to 100

Configuration Example 7: NetFlow Multiple Export Destinations

The NetFlow Multiple Export Destinations feature improves the chances of receiving complete NetFlow data by providing redundant streams of data. By sending the exact same export data to more than one NFC, fewer packets will be lost.

The NetFlow Multiple Export Destinations feature enables configuration of multiple destinations for the NetFlow data. With this feature enabled, two identical streams of NetFlow data are sent to multiple destinations. Currently, the maximum number of export destinations allowed is two. The NetFlow Multiple Export Destinations feature is only available if NetFlow is configured on networking devices supporting Cisco IOS software. The NetFlow Multiple Export Destinations feature is not supported on CatOS software.

Note Do not enter the same IP address twice. However, entering two different IP addresses with the same UDP port number is configurable.

The following configuration example shows how to configure multiple export destinations:

```
Router(config)# ip flow-export destination 10.42.42.1 9991
Router(config)# ip flow-export destination 10.0.101.254 1999
Router(config)# no ip flow-export destination 10.42.42.1 9991
```

In this example, you enable two export destinations by using the **ip flow-export destination** command twice. Using the **no ip flow-export destination** command, you disable the export of the first destination, while retaining the export of the second destination.

Configuration Example 8: NetFlow Minimum Source and Destination Mask for Prefix Aggregation Cache

The NetFlow Minimum Prefix Mask for Router-Based Aggregation feature allows you to set a minimum mask size for only an aggregation cache. The IP address that is added to the aggregation cache is appended with the maximum of the two masks: user-entered mask and the routing table mask.

The minimum mask value used by the routing device selects the granularity of the NetFlow data that will be collected:

- For coarse or low NetFlow collection granularity, select a low minimum mask value.
- For fine or high NetFlow collection granularity, select a high minimum mask value.

The mask values range from 1 to 32.

The NetFlow Minimum Prefix Mask for Router-Based Aggregation feature is only available with the router-based aggregation. Minimum masking capability is available only if router-based aggregation is enabled. This feature is only available in the following three aggregation schemes:

- Prefix aggregation
- Destination-prefix aggregation
- Source-prefix aggregation

This feature is not available in AS aggregation and Protocol-Port aggregation.

The following configuration example shows how to configure the minimum mask of a prefix aggregation cache scheme:

```
Router(config)# ip flow-aggregation cache prefix
Router(config)# mask source minimum 24
Router(config)# mask destination minimum 28
```

In this example, you configure the prefix aggregation cache using the **ip flow-aggregation cache prefix** command. Using the **mask source minimum** command, you specify the minimum value for the source mask. Using the **mask destination minimum** command, you specify the minimum value for the destination mask.

Configuration Example 9: NetFlow Minimum Destination Mask for Destination-Prefix Aggregation Cache—Configuration Example

The following configuration example shows how to configure the minimum mask of a destination-prefix aggregation cache scheme:

```
Router(config)# ip flow-aggregation cache destination-prefix
Router(config)# mask destination minimum 32
```

In this example, you configure the destination-prefix aggregation cache using the **ip flow-aggregation cache prefix** command. Using the **mask destination minimum** command, you specify the minimum value for the destination mask.

Configuration Example 10: NetFlow Multiple Export Destinations on an Aggregation Cache

The following configuration example shows how to configure multiple export destinations:

```
Router(config)# ip flow-aggregation cache destination-prefix
Router(config)# export destination 10.0.101.254 9991
Router(config)# export destination 10.0.101.254 1999
Router(config)# no export destination 10.0.101.254 1999
```

In this example, you enable a destination prefix aggregation cache using the **ip flow-aggregation cache destination-prefix** command. You enable two export destinations by using the **export destination** command twice. Using the **no export destination** command, you disable the export of the first destination, while retaining the second destination.

Configuration Example 11: Exporting a Peer AS Packet Using Version 5 Format

Figure 11 shows the flow for Exporting from a peer or Origin AS.

Figure 11 Exporting from a Peer or Origin AS

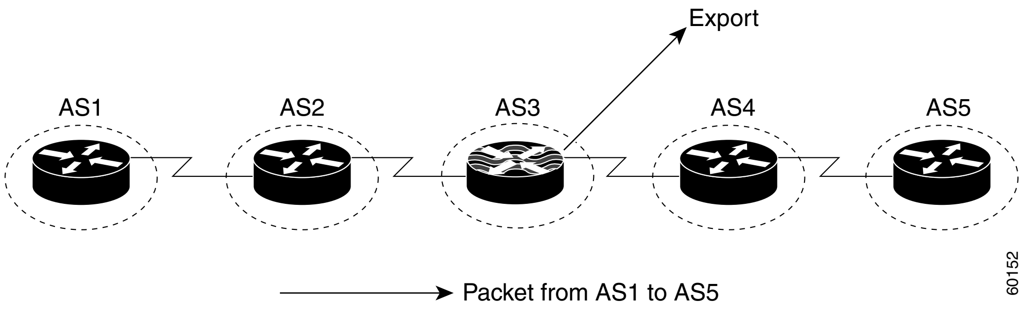

The following configuration example shows how to configure export from a peer AS using the Version 5 record format:

```
Router(config-if)# ip route-cache flow
Router(config)# ip flow-export destination 172.17.246.225 9996
Router(config)# ip flow-export version 5 peer-as
Router(config)# ip flow-export source loopback 0
Router(config)# ip flow-cache timeout
```

In this example, you configure export from a peer AS using the **ip flow-export version 5 peer-as** command. The AS source is AS2, and the AS destination is AS4.

You can also configure export from an origin AS using the **ip flow-export version 5 origin-as** command. The AS source is AS1, and the AS destination is AS5.

The AS fields will stay empty if you do not configure a peer AS or an origin AS.

Configuring MLS of NetFlow Data Export (NDE) on Catalyst Switches 65k and the 7600

NetFlow exports accounting statistics on Catalyst switches using MLS and the Version 5, Version 7 and Version 8 record formats. The following sections provide details on configuring NDE on Catalyst switches:

Configuration Example 12: Configuring the Flow Mask Mode

Configuration Example 13: Version 7 Format on the Catalyst 6500 Hybrid Mode

Configuration Example 14: Version 7 Format on the Catalyst 6500 using Native Mode RP

Configuration Example 15: Version 8 Format on the Catalyst 6500 Using CatOS

Configuration Example 16: Version 8 Format on the Catalyst 6500 Using Cisco IOS Software

Setting an MLS Flow Mask on a Catalyst Switch

Unlike Cisco routing devices, where each flow creates a NetFlow cache entry, Catalyst switches create an MLS cache. On a Catalyst switch, you can configure the flow mask of each MLS cache entry created in the following flow mask modes:

- Destination IP address only
- Source and destination IP addresses
- Full flow-mask mode (source and destination IP addresses, IP protocol, and source and destination protocol ports)

Configuration Example 12: Configuring the Flow Mask Mode

The following configuration example shows how to configure a full flow-mask mode:

```
Switch (enable)# set mls flow full
```

In this example, you specified MLS cache entries to be created with a full flow-mask mode, using the **set mls flow full** command. Cisco recommends exporting with a full flow mask.

You can also configure a default destination IP address flow mask using the **set mls flow destination** command or a default source and destination IP address flow mask using the **set mls flow destination-source** command.

Verifying MLS of NDE

The following output of the **show mls debug** command displays MLS NDE statistics:

```
Switch (enable)#  show mls debug

NDE related info:
Current Export Version : 7
Flows in nde buffer    : 3
Nde flow limit         : 27
Flow sequence          : 0
Unused flows           : 1
```

In this **show** output, a summary of MLS NDE statistics is displayed using the **show mls debug** command. The "Unused flows" output display verifies that no packets used the MLS shortcut and no records were exported.

Configuring MLS of NDE Using the Version 7 Format on Catalyst Switches

The Version 7 format uses MLS or CEF on Catalyst 6500 series switches with Supervisor 2 or above for IP unicast only.

Note Although MLS supports IP unicast, IP multicast, and IPX, the Version 7 format only supports MLS IP unicast.

Multi-layer Switching

MLS is enabled on the whole Catalyst switching device. MLS exports Layer 3 information. MLS switching is performed by the Catalyst switching device.

MLS Protocol

On a Catalyst switch, you can choose between the following operating system modes:

- Native mode—Uses CEF with Cisco IOS software.

- Hybrid mode—Uses MLS with CatOS software.

Configuration Example 13: Version 7 Format on the Catalyst 6500 Hybrid Mode

The following configuration example shows how to configure an export destination using the Version 7 record format on a Catalyst 6500:

```
Switch (enable)# set mls agingtime 256
Switch (enable)# set mls agingtime 10 5
Switch (enable)# set mls flow full
Switch (enable)# set mls nde 172.17.246.225 9996
Switch (enable)# set mls nde version 7
Switch (enable)# set mls nde enable
```

In this example, using the **set mls flow full** command will set a full flow mask.

Using the **set mls agingtime long** command, you can set the long aging timer. The default setting for the aging timer is 1,920 seconds.

On the Catalyst 6500 series switch, MLS IP is enabled by default and MLS-RP is performed internally. Note that the timers are only for the MLS cache, not the export timer.

Configuration Example 14: Version 7 Format on the Catalyst 6500 using Native Mode RP

With the Version 7 format, you typically export the Version 7 record format when the NetFlow cache is full, the flow has expired, or the timer has expired. The following configuration example shows how to configure an export destination using the Version 7 record format:

```
Router(config-if)# mls flow ip full
Router(config-if)# mls nde src_address 10.200.8.127 version 7
Router(config-if)# interface vlan1
Router(config-if)# ip route-cache flow
Router(config-if)# interface fastethernet 3/2
Router(config-if)# ip address 10.200.8.2 255.255.255.0
Router(config-if)# ip route-cache flow
Router(config-if)# ip flow-export source vlan1
Router(config-if)# ip flow-export version 7
Router(config)# ip flow-export destination 172.17.246.225 9996
```

In this example, using the **mls flow ip full** command will create a full flow mask. And setting a Version 7 export source allows you to export Version 7 to **ip flow-export destination 172.17.246.225 9996**.

Configuring MLS of NDE Using the Version 8 Format on Catalyst Switches

The Version 8 format is an enhancement that adds NetFlow support for Catalyst 5000 series switches equipped with a NetFlow feature card (NFFC) and adds support for Catalyst 6500 series switches with MSFC. The Version 8 format uses MLS or CEF on Catalyst 6500 series switches with Supervisor 2 for IP unicast only.

Configuration Example 15: Version 8 Format on the Catalyst 6500 Using CatOS

The following configuration example shows how to configure NDE using the Version 8 record format on a Catalyst 6500 with CatOS software:

```
Switch (enable)# set mls nde version 8
Switch (enable)# set mls nde 10.1.1.99 9999
Switch (enable)# set mls agingtime 8
Switch (enable)# set mls agingtime fast 8 1
Switch (enable)# set mls nde enable
```

In this example, using the set mls nde version 8 command will enable NDE using the Version 8 format.

Configuration Example 16: Version 8 Format on the Catalyst 6500 Using Cisco IOS Software

The following configuration example shows how to configure NDE using the Version 8 record format on a Catalyst 6500 with Cisco IOS software:

```
Router(config-if)# mls flow ip destination-source
Router(config-if)# mls nde src_address 10.1.1.37 version 8
Router(config-if)# interface vlan1
Router(config-if)# ip route-cache flow
Router(config-if)# interface fastethernet 3/2
Router(config-if)# ip address 10.200.8.2 255.255.255.0
Router(config-if)# ip route-cache flow
Router(config-if)# ip flow-export source vlan1
Router(config-if)# ip flow-export version 5
Router(config)# ip flow-export destination 10.1.1.99 9999
```

In this example, using the mls flow ip destination command will create a default destination IP address flow mask. And setting both a Version 8 and Version 5 export source allows you to export both Version 5 and 8 to ip flow-export destination 10.1.1.99 9999.

Related Information

To identify the hardware platform or software image information associated with a feature, use the Feature Navigator on Cisco.com to search for information about the feature or refer to the software release notes for a specific release.

NetFlow Product Marketing and Technical Documentation

http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html

The Catalyst 65k and 7600

For more information on using NetFlow on the Catalyst 65k and the 7600 see the following documents.

NetFlow IOS 12.2sx NDE configuration guide

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/122sx/swcg/nde.htm

NetFlow IOS 12.1E NDE configuration guide

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/12_1e/swconfig/nde.htm

NetFlow CatOS 8.3 NDE configuration guide

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_8_3/confg_gd/nde.htm

NetFlow on the Cisco 12000

Output NetFlow on Cisco 12000 configuration guide

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a008018883a.html

Sampled NetFlow on Cisco 12000 configuration guide

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s11/12s_sanf.htm

NetFlow on the Cisco Catalyst 4500

Catalyst 4500 12.2(20)EW configuration guide

http://www.cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a0080233fe4.html

Summary

NetFlow technology efficiently provides the metering base for a key set of applications including accounting/billing, network planning, network monitoring and outbound marketing for both service provider and enterprise customers. NetFlow does not require adoption of new or proprietary protocols or new generations of networking equipment. NetFlow is available today on most Cisco platforms. NetFlow may be deployed incrementally, on an interface-by-interface basis on strategically located edge, aggregation or WAN access routers to see a majority of IP traffic from key points in the network. NetFlow data collection and export will also serve as a key enabler for flexible, differentiated IP services based on Cisco IOS QOS. NetFlow provides an incremental path to high performance, services rich networking environments while providing maximal investment protection for the installed base of network equipment.

Appendix 1: Details Of NetFlow Export Packet Header Format For Each Export Version

Version 9

Figure 12 shows the Version 9 export packet format (table being converted to an image)

Figure 12 Version 9 Export Packet Format

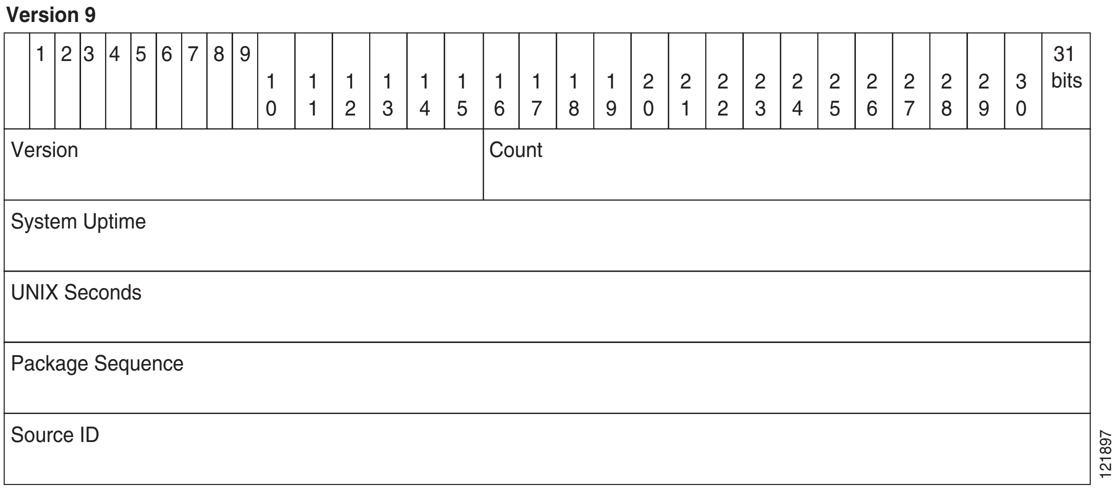

Version 9

<table>
<tr><td></td><td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td><td>9</td><td>10</td><td>11</td><td>12</td><td>13</td><td>14</td><td>15</td><td>16</td><td>17</td><td>18</td><td>19</td><td>20</td><td>21</td><td>22</td><td>23</td><td>24</td><td>25</td><td>26</td><td>27</td><td>28</td><td>29</td><td>30</td><td>31 bits</td></tr>
<tr><td colspan="16">Version</td><td colspan="16">Count</td></tr>
<tr><td colspan="32">System Uptime</td></tr>
<tr><td colspan="32">UNIX Seconds</td></tr>
<tr><td colspan="32">Package Sequence</td></tr>
<tr><td colspan="32">Source ID</td></tr>
</table>

Table 15 shows the NetFlow Version 9 Packet Header Field Descriptions.

Table 15 NetFlow Version 9 Packet Header Field Descriptions

Field Name	Value
Version	The version of NetFlow records exported in this packet; for Version 9, this value is 0x0009.
Count	Number of FlowSet records (both template and data) contained within this packet.
System Uptime	Time in milliseconds since this device was first booted.
UNIX Seconds	Seconds since 0000 Coordinated Universal Time (UTC) 1970.

Sequence Number	Incremental sequence counter of all export packets sent by this export device; this value is cumulative, and it can be used to identify whether any export packets have been missed. This is a change from the NetFlow Version 5 and Version 8 headers, where this number represented "total flows".
Source ID	The Source ID field is a 32-bit value that is used to guarantee uniqueness for all flows exported from a particular device. (The Source ID field is the equivalent of the engine type and engine ID fields found in the NetFlow Version 5 and Version 8 headers). The format of this field is vendor specific. In Cisco's implementation, the first two bytes are reserved for future expansion, and will always be zero. Byte 3 provides uniqueness with respect to the routing engine on the exporting device. Byte 4 provides uniqueness with respect to the particular line card or Versatile Interface Processor on the exporting device. Collector devices should use the combination of the source IP address plus the Source ID field to associate an incoming NetFlow export packet with a unique instance of NetFlow on a particular device.

Version 8

Figure 13 shows the version 8 export packet format

Figure 13 Version 8 Export Packet Format

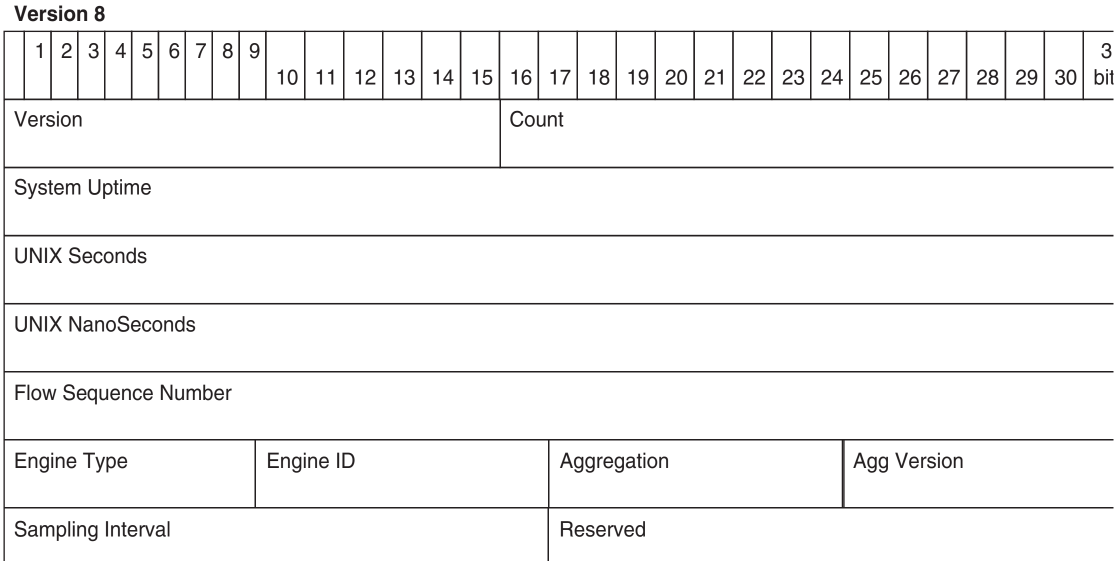

Table 16 shows the NetFlow Version 8 packet header field descriptions.

Table 16 NetFlow Version 8 Packet Header Field Descriptions

Field Name	Value
Version	The version of NetFlow records exported in this packet; for Version 7, this value is 0x0007.

Count	Number of FlowSet records (both template and data) contained within this packet.
System Uptime	Time in milliseconds since this device was first booted.
UNIX Seconds	Seconds since 0000 Coordinated Universal Time (UTC) 1970.
Unix NanoSeconds	Residual nanoseconds since 0000 UTC 1970.
Sequence Number	Sequence number of total flows seen.
Engine Type	Type of flow switching engine. Values are: • 0 for RP • 1 for VIP/LC • 2 for PFC or DFC
Engine ID	VIP or LC slot number of the flow switching engine.
aggregation	Aggregation method being used.
agg_version	Version of the aggregation export=2.

Version 7

Figure 14 shows the version 7 export packet format.

Figure 14 Version 7 Export Packet Format

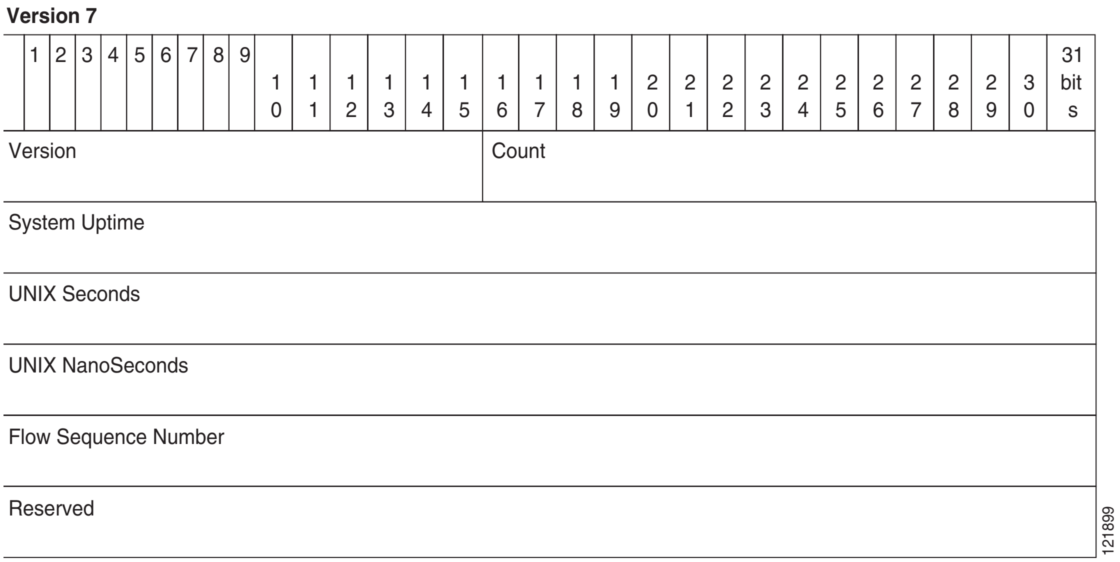

Table 17 shows the NetFlow Version 7 packet header field descriptions.

Table 17 Version 7 Packet Header Field Descriptions

Field Name	**Value**
Version	The version of NetFlow records exported in this packet; for Version 7, this value is 0x0007
Count	Number of FlowSet records (both template and data) contained within this packet
System Uptime	Time in milliseconds since this device was first booted

UNIX Seconds	Seconds since 0000 Coordinated Universal Time (UTC) 1970
Unix NanoSeconds	Residual nanoseconds since 0000 UTC 1970
Sequence Number	Sequence number of total flows seen

Figure 15 Shows the version 5 export packet format.

Figure 15 Version 5 Export Packet Format

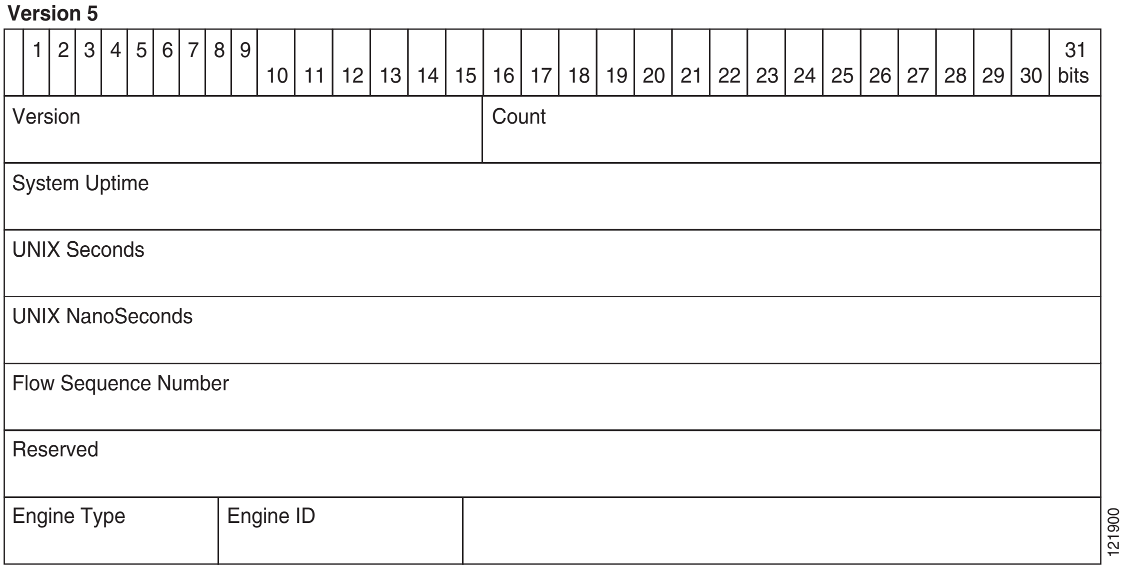

Table 18 shows the NetFlow Version 5 Packet Header Field Descriptions

Table 18 NetFlow Version 5 Packet Header Field Descriptions

Field Name	Value
Version	The version of NetFlow records exported in this packet; for Version 7, this value is 0x0007
Count	Number of FlowSet records (both template and data) contained within this packet
System Uptime	Time in milliseconds since this device was first booted
UNIX Seconds	Seconds since 0000 Coordinated Universal Time (UTC) 1970
Unix NanoSeconds	Residual nanoseconds since 0000 UTC 1970
Sequence Number	Sequence number of total flows seen
Engine Type	Type of flow switching engine, 0 for RP, 1 for VIP/LC
Engine ID	VIP or LC slot number of the flow switching engine

Version 1

Figure 16 shows the NetFlow Version 1 export packet format.

Figure 16 Version 1 Export Packet Format

Table 19 shows the NetFlow Version 1 packet header field descriptions.

Table 19 NetFlow Version 1 Packet Header Field Descriptions

Field Name	Value
Version	The version of NetFlow records exported in this packet; for Version 7, this value is 0x0007
Count	Number of FlowSet records (both template and data) contained within this packet
System Uptime	Time in milliseconds since this device was first booted
UNIX Seconds	Seconds since 0000 Coordinated Universal Time (UTC) 1970
Unix NanoSeconds	Residual nanoseconds since 0000 UTC 1970

Appendix 2: Details for NetFlow Export Formats

Table 20 shows a typical version 9 export format.

Table 20 A typical Version 9 Export Format

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
FlowSet ID = 0															
Length															
Template ID															
Field Count															
Field 1 Type															
Field 1 Length															
Field 2 Type															

Field 2 Length
—
Field *N* Type
Field *N* Length
Template ID
Field Count
Field 1 Type
Field 1 Length
Field 2 Type
Field 2 Length
—
Field *N* Type
Field *N* Length

Table 21 shows the NetFlow Version 9 template flowset field descriptions.

Table 21 NetFlow Version 9 Template FlowSet Field Descriptions

Field Name	**Value**
FlowSet ID	The FlowSet ID is used to distinguish template records from data records. A template record always has a FlowSet ID in the range of 0-255. Currently template record that describes flow fields has a FlowSet ID of zero and the template record that describes option fields (described below) has a FlowSet ID of 1. A data record always has a nonzero FlowSet ID greater than 255.
Length	Length refers to the total length of this FlowSet. Because an individual template FlowSet may contain multiple template IDs (as illustrated above), the length value should be used to determine the position of the next FlowSet record, which could be either a template or a data FlowSet. Length is expressed in type/length/value (TLV) format, meaning that the value includes the bytes used for the FlowSet ID and the length bytes themselves, as well as the combined lengths of all template records included in this FlowSet.
Template ID	As a router generates different template FlowSets to match the type of NetFlow data it will be exporting, each template is given a unique ID. This uniqueness is local to the router that generated the template ID. Templates that define data record formats begin numbering at 256 since 0-255 are reserved for FlowSet IDs.

Field Count	This field gives the number of fields in this template record. Because a template FlowSet may contain multiple template records, this field allows the parser to determine the end of the current template record and the start of the next.
Field Type	This numeric value represents the type of the field. The possible values of the field type are vendor specific. Cisco supplied values are consistent across all platforms that support NetFlow Version 9. At the time of the initial release of the NetFlow Version 9 code (and after any subsequent changes that could add new field-type definitions), Cisco provides a file that defines the known field types and their lengths. The currently defined field types are detailed in Table 6.
Field Length	This number gives the length of the above-defined field, in bytes.

More information about the fields currently available in version 9 and the export format architecture are available in the NetFlow Version 9 Flow-Record Format document**.**

Table 22 shows the NetFlow Version 1 export packet format.

Table 22 NetFlow Version 1 Export Packet Format

Content	Bytes	Description
srcaddr	0-3	Source IP address
dstaddr	4-7	Destination IP address
nexthop	8-11	Next hop router's IP address
input	12-13	Ingress interface SNMP ifIndex
output	14-15	Egress interface SNMP ifIndex
dPkts	16-19	Packets in the flow
dOctets	20-23	Octets (bytes) in the flow
first	24-27	SysUptime at start of the flow
last	28-31	SysUptime at the time the last packet of the flow was received
srcport	32-33	Layer 4 source port number or equivalent
dstport	34-35	Layer 4 destination port number or equivalent
pad1	36	Unused (zero) byte
prot	37	Layer 4 protocol (for example, 6=TCP, 17=UDP)
tos	38	IP type-of-service byte
tcp_flags	39	Cumulative OR of TCP flags
Pad 2	40	Pad 2 is unused (zero) bytes
Pad 3	41-42	Pad 3 is unused (zero) bytes
reserved	43	Autonomous system number of the destination, either origin or peer

Catalyst 65k NDE Versions

NDE on the PFC supports the following NDE versions to export the statistics captured on the PFC for Layer 3-switched traffic:

- Supervisor Engine 1 and PFC—NDE version 7
- Supervisor Engine 2 and PFC2
  - NDE version 5 with Release 12.1(13)E and later releases
  - NDE version 7 with all releases

Depending on the current flow mask, some fields in the flow records might not have values. When the PFC exports cached entries, unsupported fields are filled with a zero (0).

The following tables list the supported NDE fields:

- Table 23—Version 5 header format
- Table 24—Version 5 flow record format
- Table 25—Version 7 header format
- Table 26—Version 7 flow record format

Table 23 NDE Version 5 Header Format

Bytes	Content	Description
0–1	version	Netflow export format version number
2–3	count	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
20–21	engine_type	Type of flow switching engine
21–23	engine_id	Slot number of the flow switching engine

Table 24 NDE Version 5 Flow Record Format

Bytes	Content	Description	**Flow masks:** **• X=Populated** **• A=Additional field** (see the "Populating Additional NDE Fields" section on page 49)				
			Destination	**Destination Source**	**Destination Source Interface**[1]	**Full**	**Full Interface**[1]
0–3	srcaddr	Source IP address		X	X	X	X
4–7	dstaddr	Destination IP address	X	X	X	X	X

Table 24 NDE Version 5 Flow Record Format (continued)

			Flow masks: • **X=Populated** • **A=Additional field** (see the "Populating Additional NDE Fields" section on page 49)				
Bytes	**Content**	**Description**	**Destination**	**Destination Source**	**Destination Source Interface**[1]	**Full**	**Full Interface**[1]
8–11	nexthop	Next hop router's IP address	**A**[2]	**A**	**A**	**A**	**A**
12–13	input	Ingress interface SNMP ifIndex			**X**		**X**
14–15	output	Egress interface SNMP ifIndex	**A**[2]	**A**	**A**	**A**	**A**
16–19	dPkts	Packets in the flow	**X**	**X**	**X**	**X**	**X**
20–23	dOctets	Octets (bytes) in the flow	**X**	**X**	**X**	**X**	**X**
24–27	first	SysUptime at start of the flow	**X**	**X**	**X**	**X**	**X**
28–31	last	SysUptime at the time the last packet of the flo w was received	**X**	**X**	**X**	**X**	**X**
32–33	srcport	Layer 4 source port number or equivalent				**X**	**X**
34–35	dstport	Layer 4 destination port number or equivalent				**X**	**X**
36	pad1	Unused (zero) byte					
37	tcp_flags	Cumulative OR of TCP flags					
38	prot	Layer 4 protocol (for example, 6=TC P, 17=UDP)				**X**	**X**
39	tos	IP type-of-service byte					
40–41	src_as	Autonomous system number of the source, either origin or peer		**A**	**A**	**A**	**A**
42–43	dst_as	Autonomous system number of the destination, either origin or peer	**A**	**A**	**A**	**A**	**A**

Table 24 NDE Version 5 Flow Record Format (continued)

			Flow masks: **• X=Populated** **• A=Additional field** (see the "Populating Additional NDE Fields" section on page 49)				
Bytes	**Content**	**Description**	**Destination**	**Destination Source**	**Destination Source Interface**[1]	**Full**	**Full Interface**[1]
44–45	src_mask	Source address prefix mask bits					
46–47	dst_mask	Destination address prefix mask bits					
48	pad2	Pad 2 is unused (zero) bytes					

[1] Supported in Release 12.1(13)E and later releases.

[2] With the destination flow-mask, the "Next hop router's IP address" field and the "Output interface's SNMP ifIndex" field might not contain information that is accurate for all flows.

Table 25 NDE Version 7 Header Format

Bytes	**Content**	**Description**
0–1	version	Netflow export format version number
2–3	count	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
20–24	reserved	Unused (zero) bytes

Table 26 NDE Version 7 Flow Record Format

			Flow masks: **• X=Populated** **• A=Additional field** (see the "Populating Additional NDE Fields" section on page 49)				
Bytes	**Content**	**Description**	**Destination**	**Destination Source**	**Destination Source Interface**[1]	**Full**	**Full Interface**[1]
0–3	srcaddr	Source IP address		**X**	**X**	**X**	**X**
4–7	dstaddr	Destination IP address	**X**	**X**	**X**	**X**	**X**
8–11	nexthop	Next hop router's IP address	**X**[2]	**X**	**X**	**X**	**X**

Table 26 NDE Version 7 Flow Record Format (continued)

			Flow masks: • **X=Populated** • **A=Additional field** (see the "Populating Additional NDE Fields" section on page 49)				
Bytes	**Content**	**Description**	**Destination**	**Destination Source**	**Destination Source Interface**[1]	**Full**	**Full Interface**[1]
12–13	input	Ingress interface SNMP ifIndex			X		X
14–15	output	Egress interface SNMP ifIndex	X[2]	X	X	X	X
16–19	dPkts	Packets in the flow	X	X	X	X	X
20–23	dOctets	Octets (bytes) in the flow	X	X	X	X	X
24–27	First	SysUptime at start of the flow	X	X	X	X	X
28–31	Last	SysUptime at the time the last packet of the flow was received	X	X	X	X	X
32–33	srcport	Layer 4 source port number or equivalent				X	X
34–35	dstport	Layer 4 destination port number or equivalent				X	X
36	flags	flow mask in use	X	X	X	X	X
37	tcp_flags	Cumulative OR of TCP flags					
38	prot	Layer 4 protocol (for example, 6=TCP, 17=UDP)				X	X
39	tos	IP type-of-service byte					
40–41	src_as	Autonomous system number of the source, either origin or peer		A	A	A	A
42–43	dst_as	Autonomous system number of the destination, either origin or peer	A	A	A	A	A
44–45	src_mask	Source address prefix mask bits					
46–47	dst_mask	Destination address prefix mask bits					

Table 26 NDE Version 7 Flow Record Format (continued)

Bytes	Content	Description	**Flow masks:** • **X=Populated** • **A=Additional field** (see the “Populating Additional NDE Fields” section on page 49)				
			Destination	**Destination Source**	**Destination Source Interface**[1]	**Full**	**Full Interface**[1]
48	pad2	Pad 2 is unused (zero) bytes					
49–50	MLS RP	IP address of MLS router	X	X	X	X	X

[1] Supported in Release 12.1(13)E and later releases.

[2] With the destination flow-mask, the “Next hop router’s IP address” field and the “Output interface’s SNMP ifIndex” field might not contain information that is accurate for all flows.

Populating Additional NDE Fields

With Release 12.1(13)E and later releases, you can configure NDE to populate the following additional fields in the NDE packets:

- IP address of the next hop router
- Egress interface SNMP ifIndex
- Source autonomous system number
- Destination autonomous system number

Not all of the additional fields are populated with all flow masks. See the “Catalyst 65k NDE Versions” section on page 45 for additional information.

To populate the additional fields in NDE packets, perform this task:

Command	Purpose
`Router(config)# mls nde interface`	Populates additional fields in NDE packets.
`Router(config)# no mls nde interface`	Disables population of the additional fields.

This example shows how to populate the additional fields in NDE packets:

```
Router(config)# mls nde interface
```

Appendix 3: Router Based Aggregation Schemes And Detailed NetFlow Export Formats

Selecting a NetFlow Aggregation Cache Scheme

You can configure each aggregation cache scheme with its individual cache size, cache ager timeout parameter, export destination IP address, and export destination UDP port. As data flows expire in the main cache (depending on the aggregation scheme configured), relevant information is extracted from the expired flow and the corresponding flow entry in the aggregation cache is updated. Each aggregation cache contains different field combinations that determine which data flows are grouped. The default aggregation cache size is 4096. The NetFlow aggregation cache schemes are described in the following sections:

- AS Aggregation Scheme
- Destination-Prefix Aggregation Scheme
- Prefix Aggregation Scheme
- Protocol-Port Aggregation Scheme
- Source Prefix Aggregation Scheme

The NetFlow ToS-Based Router Aggregation feature introduces support for six aggregation cache schemes that include the ToS byte as a field. The NetFlow ToS-Based Router Aggregation feature provides the ability to enable limited router-based ToS aggregation of NetFlow data, which results in summarized NetFlow data being exported to a collection device. The result is lower bandwidth requirements for NetFlow data and reduced platform requirements for NetFlow data collection devices.

This support is described in the following sections:

- AS-ToS Aggregation Scheme
- Destination-Prefix-ToS Aggregation Scheme
- Prefix-ToS Aggregation Scheme
- Protocol-Port-ToS Aggregation Scheme
- Source Prefix-ToS Aggregation Scheme
- Prefix-Port Aggregation Scheme

AS Aggregation Scheme

The AS aggregation scheme provides substantial NetFlow export data volume reduction and generates AS-to-AS traffic flow data. The scheme groups data flows with the same source BGP AS, destination BGP AS, input interface, and output interface. See Figure 17 for the AS aggregation data export format.

The aggregated NDE records report the following:

- Source and destination BGP AS
- Number of flows summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Output and input interfaces
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 17 ***AS Aggregation Data Export Format***

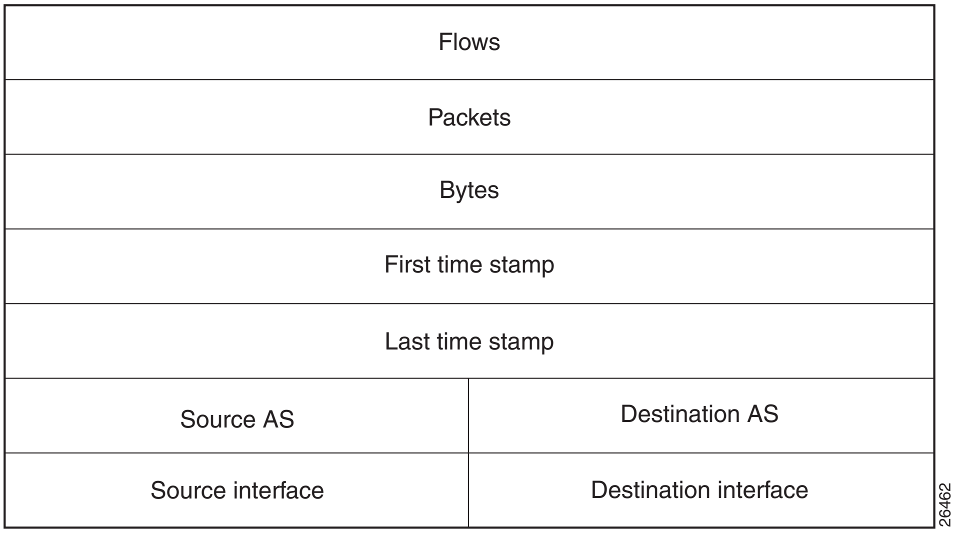

Flows	
Packets	
Bytes	
First time stamp	
Last time stamp	
Source AS	Destination AS
Source interface	Destination interface

Table 27 describes the bytes and data fields used in the AS aggregation export record format.

Table 27 **AS Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 21	src_as	AS of the source IP address (peer or origin).
22 to 23	dst_as	AS of the destination IP address (peer or origin).
24 to 25	input	SNMP index of the input interface.
26 to 27	output	SNMP index of the output interface.

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 AS, and output interface. See Figure 18 for the Destination-Prefix aggregation data export format.

The aggregated NDE records report the following:

- Destination prefix
- Destination prefix mask
- Destination BGP AS
- Number of flows summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Output interface
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 18 Destination-Prefix Aggregation Data Export Format

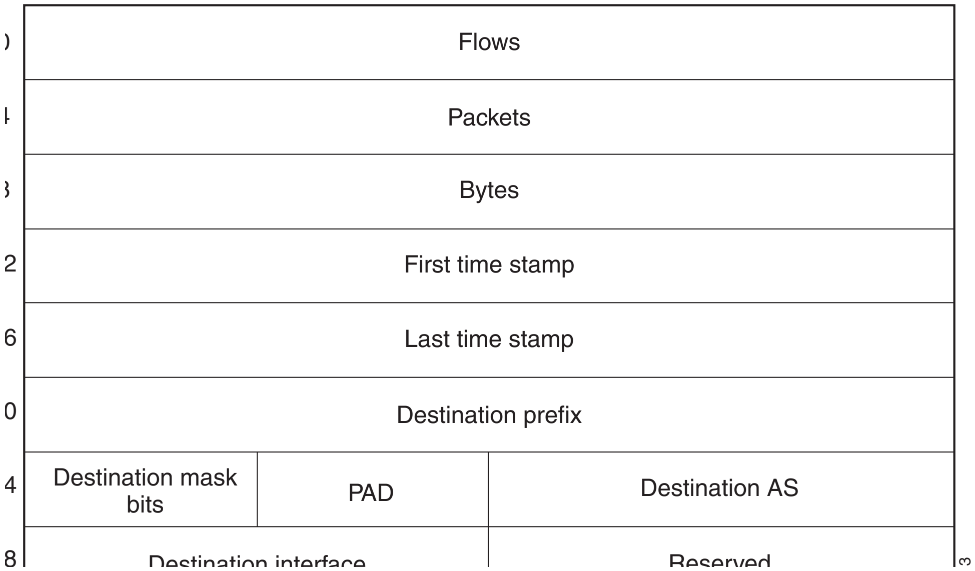

Table 28 describes the bytes and data fields used in the Destination-Prefix aggregation export record format.

Table 28 **Destination-Prefix Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.

Bytes	Content	Definition
20 to 23	dst_prefix	Prefix that the destination IP address of the aggregated flows belonged to.
24	dst_mask	Number of bits in the destination address prefix mask.
25	pad	Zero field.
26 to 27	dst_as	AS of the destination IP address (peer or origin).
28 to 29	output	SNMP index of the output interface.
30 to 31	reserved	Zero field.

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 AS, destination BGP AS, input interface, and output interface. See Figure 19 for the Prefix Aggregation data export format:

The aggregated NDE records report the following:

- Source and destination prefix
- Source and destination prefix mask
- Source and destination BGP AS
- Number of flows summarized by the aggregated record
- Number of packets summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Input and output interface
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 19 Prefix Aggregation Data Export Format

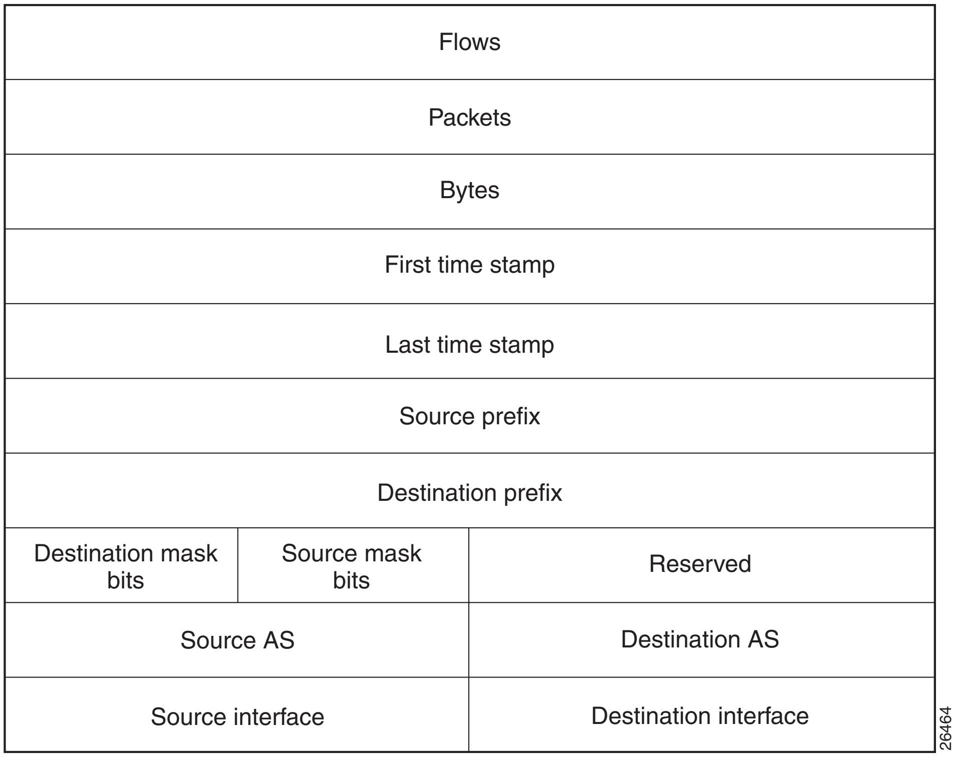

Flows		
Packets		
Bytes		
First time stamp		
Last time stamp		
Source prefix		
Destination prefix		
Destination mask bits	Source mask bits	Reserved
Source AS		Destination AS
Source interface		Destination interface

Table 29describes the bytes and data fields used in the Prefix aggregation export record format.

Table 29 **Prefix Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.

Bytes	Content	Definition
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 23	src_prefix	Prefix that the source IP address of the aggregated flows belonged to.
24 to 27	dst_prefix	Prefix that the destination IP address of the aggregated flows belonged to.
28	dst_mask	Number of bits in the destination address prefix mask.
29	src_mask	Number of bits in the source address prefix mask.
30	pad	Zero field.
31 to 32	src_as	AS of the source IP address (peer or origin).
33 to 34	dst_as	AS of the destination IP address (peer or origin).
35 to 36	input	SNMP index of the input interface.
37 to 38	output	SNMP index of the output interface.

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 20 for the Protocol-Port aggregation data export format.

The aggregated NDE records report the following:

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

Figure 20 Protocol-Port Aggregation Data Export Format

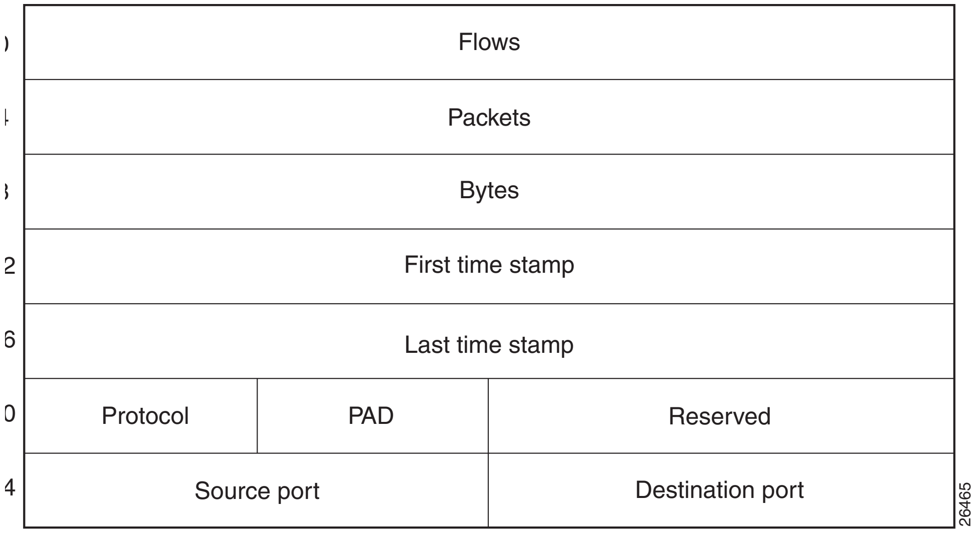

Table 30 describes the bytes and data fields used in the Protocol-Port aggregation export record format.

Table 30 **Protocol-Port Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20	prot	IP protocol byte.
21	pad	Zero field.
22 to 23	reserved	Zero field.

Bytes	Content	Definition
24 to 25	srcport	Source UDP or TCP port number.
26 to 27	dstport	Destination UDP or TCP port number.

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 AS, and input interface. See Figure 21 for the Source Prefix aggregation data export format.

The aggregated NDE records report the following:

- Source prefix
- Source prefix mask
- Source BGP AS
- Number of packets summarized by the aggregated record
- Number of bytes summarized by the aggregated record
- Input interface
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 21 Source Prefix Aggregation Data Export Format

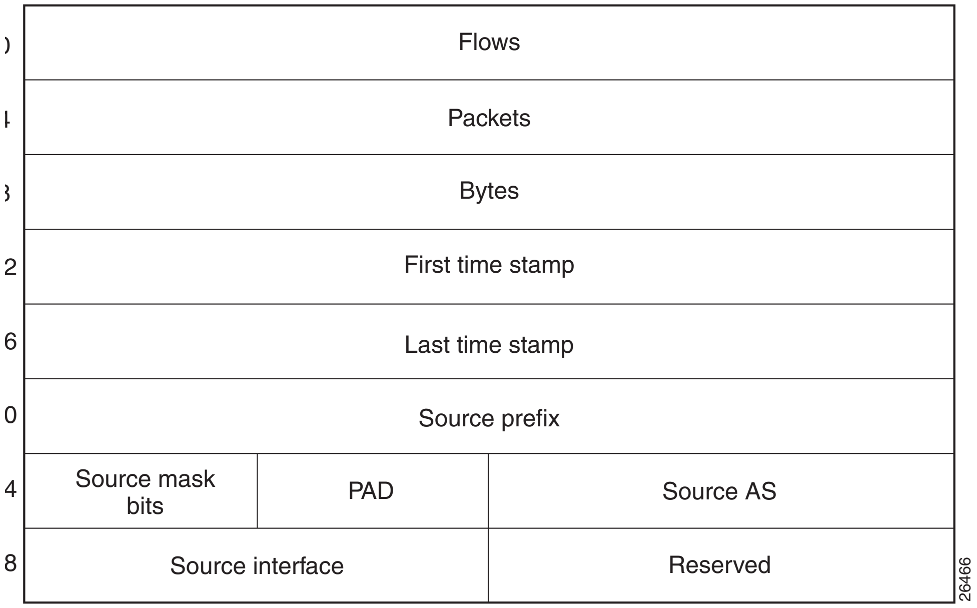

Table 31 describes the bytes and data fields used in the Source Prefix aggregation export record format.

Table 31 Source Prefix Aggregation Export Record Format

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.

Bytes	Content	Definition
20 to 23	src_prefix	Prefix that the source IP address of the aggregated flows belonged to.
24	src_mask	Number of bits in the source address prefix mask.
25	pad	Zero field.
26 to 27	src_as	AS of the source IP address (peer or origin).
28 to 29	input	SNMP index of the input interface.
30 to 31	reserved	Zero field.

AS-ToS Aggregation Scheme

The AS-ToS aggregation scheme groups together flows that have the same source and destination BGP AS, source and destination interfaces, and ToS byte. This aggregation scheme is particularly useful for generating AS-to-AS traffic flow data, and for providing substantial NetFlow export data volume reduction. See Figure 22 for the AS-ToS aggregation export record format.

The aggregated NDE records report the following:

- Source BGP AS
- Destination BGP AS
- ToS byte
- Number of flows summarized by the aggregated record
- Number of bytes summarized by this aggregated record
- Number of packets summarized by this aggregation record
- Source and destination interface
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 22 AS-ToS Aggregation Export Record Format

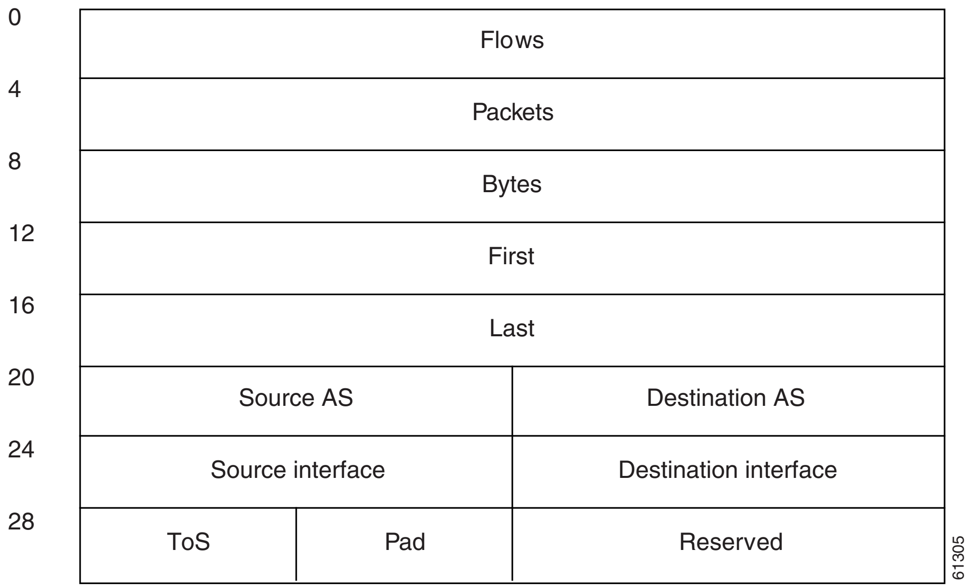

Table 32 describes the bytes and data fields used in the AS-ToS aggregation export record format.

Table 32 **AS-ToS Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 21	src_as	AS of the source IP address (peer or origin).

Bytes	Content	Definition
22 to 23	dst_as	AS of the destination IP address (peer or origin).
24 to 25	input	SNMP index of the input interface.
26 to 27	output	SNMP index of the output interface.
28	tos	ToS byte.
29	pad	Zero field.
30 to 31	reserved	Zero field.

Destination-Prefix-ToS Aggregation Scheme

The Destination-Prefix-ToS aggregation scheme groups flows with common destination prefix, destination prefix mask, destination BGP AS, ToS byte, and output interface. This aggregation scheme is particularly useful for generating data with which to examine the destinations of network traffic passing through a NetFlow-enabled device. See Figure 23 for the Destination-Prefix aggregation export record format.

The aggregated NDE records report the following:

- Destination IP address
- Destination prefix mask
- Destination AS
- ToS byte
- Number of flows summarized by the aggregated record
- Number of bytes summarized by this aggregated record
- Number of packets summarized by this aggregation record
- Output interface
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 23 Destination-Prefix Aggregation Export Record Format

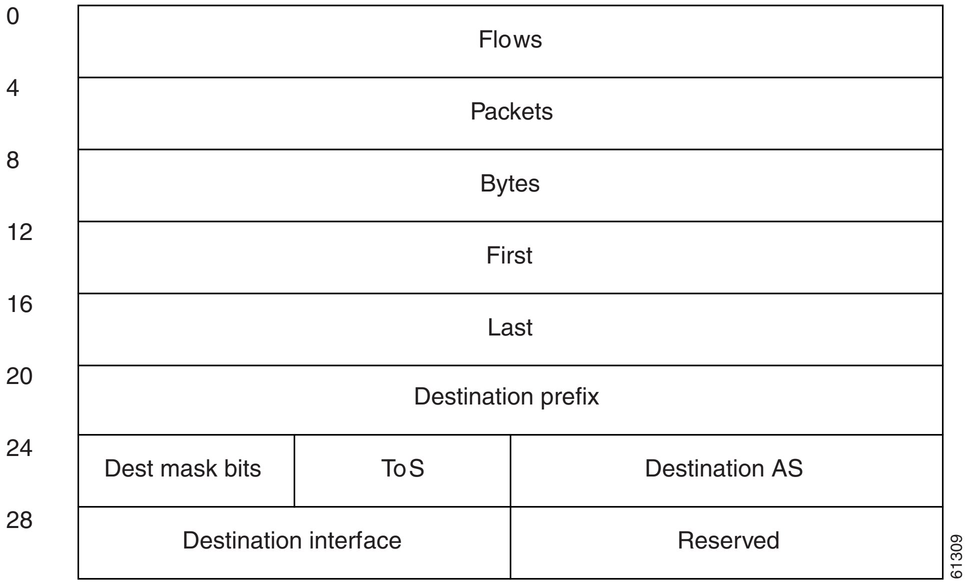

0	Flows		
4	Packets		
8	Bytes		
12	First		
16	Last		
20	Destination prefix		
24	Dest mask bits	ToS	Destination AS
28	Destination interface		Reserved

Table 33 describes the bytes and data fields used in the Destination-Prefix aggregation export record format.

Table 33 **Destination-Prefix Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.

Bytes	Content	Definition
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 23	dst_prefix	Prefix that the destination IP address of the aggregated flows belonged to.
24	dst_mask	Number of bits in the destination address prefix mask.
25	tos	ToS byte.
26 to 27	dst_as	AS of the destination IP address (peer or origin).
28 to 29	output	SNMP index of the output interface.
30 to 31	reserved	Zero field.

Prefix-ToS Aggregation Scheme

The Prefix-ToS aggregation scheme groups together flows with common source prefix, source mask, destination prefix, destination mask, source BGP AS, destination BGP AS, input interface, output interface, and ToS byte. This aggregation scheme is particularly useful for generating data with which to examine the sources and destinations of network traffic passing through a NetFlow-enabled device. See Figure 24 for the Prefix-ToS aggregation export record format.

The aggregated NDE records report the following:

- Source prefix
- Source prefix mask
- Destination prefix
- Destination prefix mask
- Source AS
- Destination AS
- Source interface
- Destination interface
- ToS byte
- Number of flows summarized by the aggregated record
- Number of bytes summarized by this aggregated record
- Number of packets summarized by this aggregation record
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 24 Prefix-ToS Aggregation Export Record Format

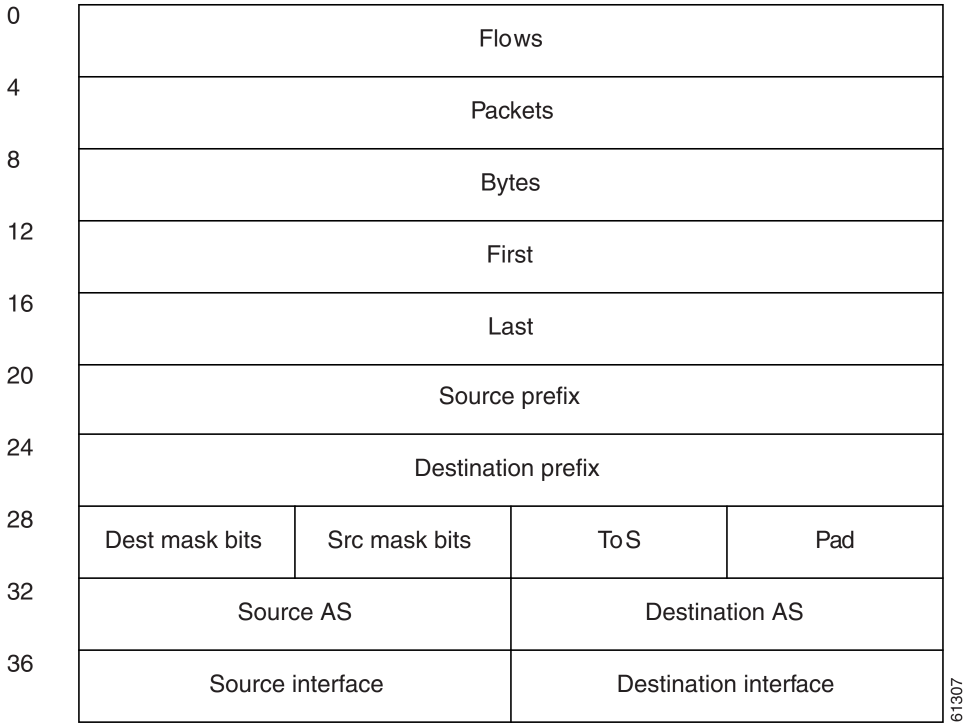

Table 34 describes the bytes and data fields used in the Prefix-ToS aggregation export record format.

Table 34 **Prefix-ToS Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 23	src_prefix	Prefix that the source IP address of the aggregated flows belonged to.
24 to 27	dst_prefix	Prefix that the destination IP address of the aggregated flows belonged to.
28	dst_mask	Number of bits in the destination address prefix mask.
29	src_mask	Number of bits in the source address prefix mask.
30	tos	ToS byte.
31	pad	Zero field.
32 to 33	src_as	AS of the source IP address (peer or origin).
34 to 35	dst_as	AS of the destination IP address (peer or origin).
36 to 37	input	SNMP index of the input interface.
38 to 39	output	SNMP index of the output interface.

Protocol-Port-ToS Aggregation Scheme

The Protocol-Port-ToS aggregation scheme groups flows with common IP protocol, ToS byte, source and destination port numbers when applicable, and source and destination interfaces. This aggregation scheme is particularly useful for generating data with which to examine network usage by type of traffic. See Figure 25 for the Protocol-Port-ToS aggregation export record format.

The aggregated NetFlow export record reports the following:

- Source application port number
- Destination port number
- Source and destination interface
- IP protocol
- ToS byte
- Number of flows summarized by the aggregated record
- Number of bytes summarized by this aggregated record
- Number of packets summarized by this aggregation record
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 25 ***Protocol-Port-ToS Aggregation Export Record Format***

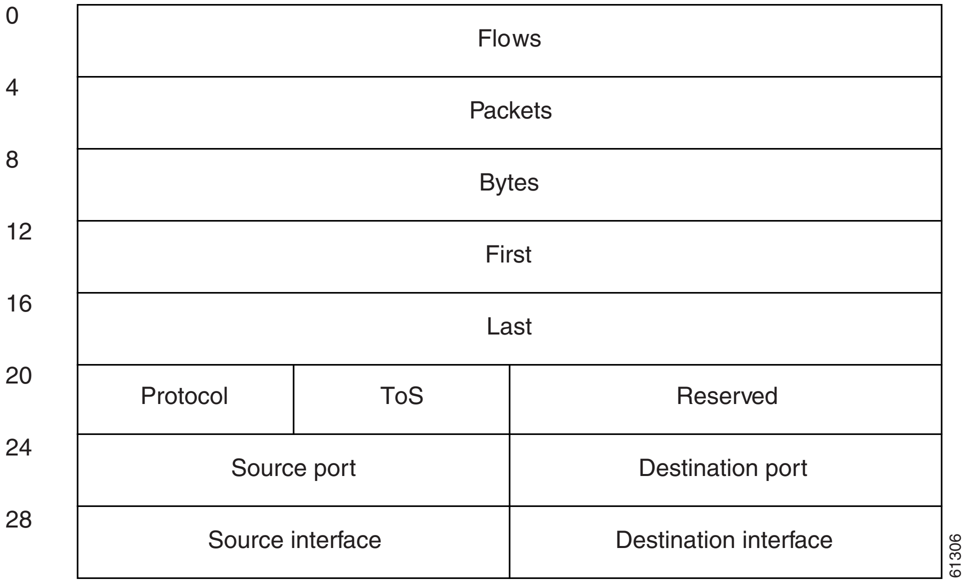

Table 35 describes the bytes and data fields used in the Protocol-Port-ToS aggregation export record format.

Table 35 **Protocol-Port-ToS Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.

Bytes	Content	Definition
16 to 19	Last	SysUptime at which the last packet was switched.
20	prot	IP protocol byte.
21	tos	ToS byte.
22 to 23	reserved	Zero field.
24 to 25	srcport	Source UDP or TCP port number.
26 to 27	dstport	Destination UDP or TCP port number.
28 to 29	input	SNMP index of the interface the packets arrived on.
30 to 31	output	SNMP index of the interface the packets went out.

Source Prefix-ToS Aggregation Scheme

The Source Prefix-ToS aggregation scheme groups flows with common source prefix, source prefix mask, source BGP AS, ToS byte, and input interface. This aggregation scheme is particularly useful for generating data with which to examine the sources of network traffic passing through a NetFlow-enabled device. See Figure 26 for the Source Prefix-ToS aggregation export record format.

The aggregated NDE records report the following:

- Source prefix
- Source prefix mask
- Source AS
- ToS byte
- Number of bytes summarized by this aggregated record
- Number of packets summarized by this aggregation record
- Input interface
- Time stamps of when the first packet is switched and when the last packet is switched.

Note When a routing device does not have a prefix for the source IP address in the flow, 0.0.0.0 with 0 mask bits is used rather than making /32 entries to prevent DOS attacks with random source address from thrashing the aggregation caches. This is done for the destination in the Destination-Prefix-ToS, and the Prefix-ToS and Prefix-Port aggregation schemes.

Figure 26 Source Prefix-ToS Aggregation Export Record Format

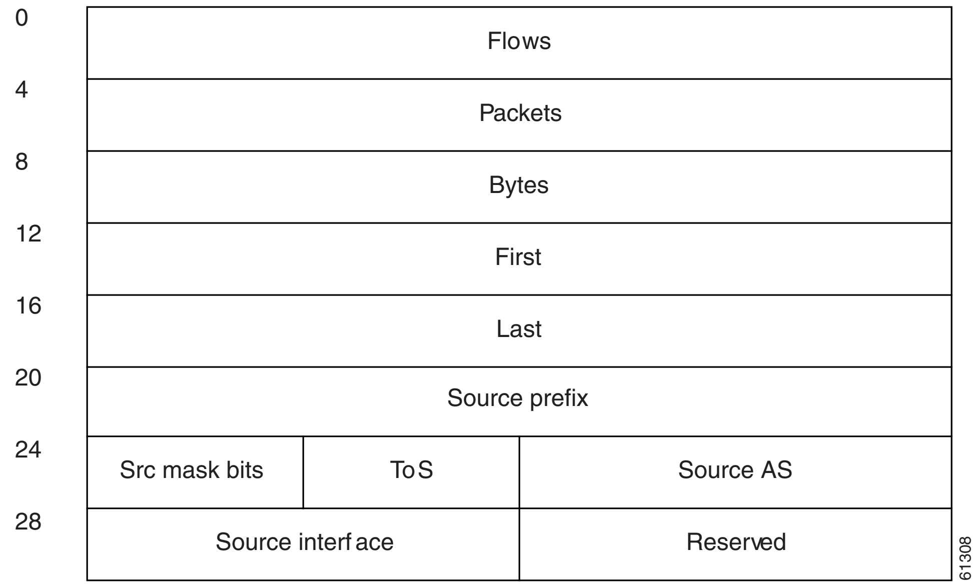

Table 36 describes the bytes and data fields used in the Source Prefix-ToS aggregation export record format.

Table 36 **Source Prefix-ToS Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 23	src_prefix	Prefix that the source IP address of the aggregated flows belonged to.
24	src_mask	Number of bits in the source address prefix mask.
25	tos	ToS byte.
26 to 27	src_as	AS of the source IP address (peer or origin).
28 to 29	input	SNMP index of the input interface.
30 to 31	reserved	Zero field.

Prefix-Port Aggregation Scheme

The Prefix-Port aggregation scheme groups flows with common source prefix, source mask, destination prefix, destination mask, source port and destination port when applicable, input interface, output interface, protocol, and ToS byte. This aggregation scheme is particularly useful for generating data with which to examine the sources and destinations of network traffic passing through a NetFlow-enabled device. See Figure 27 for the Prefix-Port aggregation export record format.

The aggregated NDE records report the following:

- Source prefix
- Source prefix mask
- Destination prefix
- Destination prefix mask
- Source port
- Destination port
- Source interface
- Destination interface
- Protocol
- ToS byte
- Number of flows summarized by the aggregated record
- Number of bytes summarized by this aggregated record
- Number of packets summarized by this aggregation record
- Time stamps of when the first packet is switched and when the last packet is switched

Figure 27 ***Prefix-Port Aggregation Export Record Format***

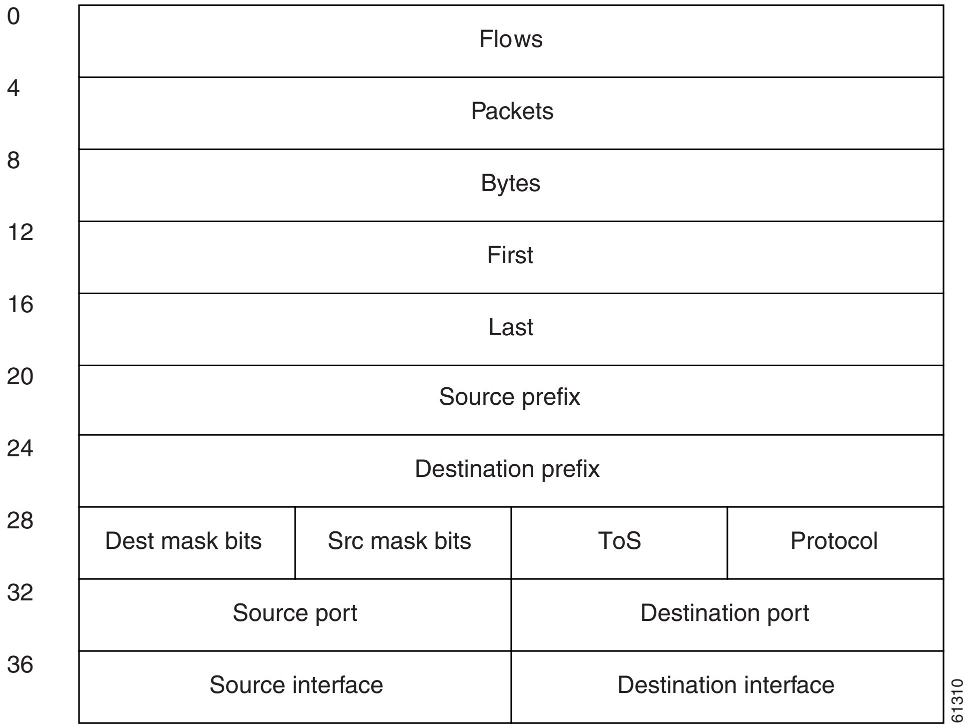

Table 37 describes the bytes and data fields used in the Prefix-Port aggregation export record format.

Table 37 **Prefix-Port Aggregation Export Record Format**

Bytes	Content	Definition
0 to 3	flows	Number of main cache flows that were aggregated.
4 to 7	dPkts	Number of packets in the aggregated flows.
8 to 11	dOctets	Number of bytes in the aggregated flows.
12 to 15	First	SysUptime at which the first packet was switched.
16 to 19	Last	SysUptime at which the last packet was switched.
20 to 23	src_prefix	Prefix that the source IP address of the aggregated flows belonged to.
24 to 27	dst_prefix	Prefix that the destination IP address of the aggregated flows belonged to.
28	dst_mask	Number of bits in the destination prefix.
29	src_mask	Number of bits in the source prefix.
30	tos	ToS byte.
31	prot	IP protocol byte.
32 to 33	srcport	Source UDP or TCP port number.
34 to 35	dstport	Destination UDP or TCP port number.
36 to 37	input	SNMP index of the input interface.
38 to 39	output	SNMP index of the output interface.