Configuring NetFlow and NDE

This chapter describes how to configure NetFlow statistics collection and NetFlow Data Export (NDE) on the Cisco 7600 series routers.

Note

- For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS NetFlow Command Reference* at this URL:

  http://www.cisco.com/en/US/docs/ios/netflow/command/reference/nf_book.html
- NetFlow version 9 is supported—See this publication:

  Cisco IOS NetFlow Configuration Guide
- NetFlow multicast support includes the NetFlow v9 export format feature. See this publication:

  Cisco IOS NetFlow Configuration Guide

  You do not need to configure multicast fast switching or multicast distributed fast switching (MDFS); multicast CEF switching is supported.

This chapter consists of these sections:

Understanding How NetFlow and NDE Work

These sections describe how NetFlow and NDE work:

- NetFlow and NDE on the PFC, page 56-2

NetFlow and NDE Overview

NetFlow collects statistics about traffic that flows through the router. NetFlow Data Export (NDE) enables you to export those statistics to an external data collector for analysis.

NetFlow and NDE are either enabled globally or enabled on individual interfaces, depending on which software release you are using:

- In Cisco IOS Release 12.2SRA and earlier releases, NetFlow is enabled globally, which means that statistics are gathered for all interfaces on the router.
- In Cisco IOS Release 12.2SRB, you can enable NetFlow on individual interfaces for IPv4 traffic on Layer 3 interfaces. NetFlow for IPv6 traffic continues to operate in global mode. For more information about this feature, see the "Per-Interface NetFlow and NDE" section on page 56-10.

Note Beginning in Release 12.2SRB, global-mode NetFlow for IPv4 traffic is no longer the default. To achieve the same global-mode functionality as before, you must now manually enable NetFlow on each Layer 3 interface where you want to capture statistics for IPv4 traffic flows.

You can collect statistics for both routed and bridged traffic. Note, however, that the PFC3A collects statistics only for routed traffic.

You can configure two external data collector addresses, which improves the probability of receiving complete NetFlow data by providing redundant data streams with a PFC3.

To reduce the volume of statistics collected, use:

- NetFlow Sampling, which reduces the number of statistics collected
- NetFlow aggregation, which merges collected statistics

NetFlow and NDE on the MSFC

The NetFlow cache on the MSFC captures statistics for flows routed in software. The MSFC supports NetFlow aggregation for traffic routed in software. For more information, see the Cisco IOS NetFlow Configuration Guide.

The MSFC supports NetFlow ToS-based router aggregation, For more information, see the Cisco IOS NetFlow Configuration Guide.

NetFlow and NDE on the PFC

The NetFlow cache on the PFC captures statistics for flows routed in hardware. The PFC supports sampled NetFlow and NetFlow aggregation for traffic routed in hardware. The PFC does not support NetFlow ToS-Based Router Aggregation.

These sections describe NetFlow and NDE on the PFC in more detail:

- Flow Masks, page 56-3
- NDE Versions, page 56-3
- MLS Cache Entries, page 56-7

- NetFlow Sampling, page 56-7
- NetFlow Aggregation, page 56-10

Flow Masks

This section describes the flow masks that are used to create NetFlow entries. Two sets of flow masks are available: for Release 12.2SRA and Release 12.2SRB. NetFlow applies the selected flow mask to all statistics gathered on the router.

Release 12.2SRA

Cisco IOS Release 12.2SRA uses the following types of flow masks to create NetFlow entries:

- source-only—A less-specific flow mask. The PFC maintains one entry for each source IP address. All flows from a given source IP address use this entry.
- destination—A less-specific flow mask. The PFC maintains one entry for each destination IP address. All flows to a given destination IP address use this entry.
- destination-source—A more-specific flow mask. The PFC maintains one entry for each source and destination IP address pair. All flows between same source and destination IP addresses use this entry.
- destination-source-interface—A more-specific flow mask. Adds the source VLAN SNMP ifIndex to the information in the destination-source flow mask.
- full—A more-specific flow mask. The PFC creates and maintains a separate cache entry for each IP flow. A full entry includes the source IP address, destination IP address, protocol, and protocol interfaces.
- full-interface—The most-specific flow mask. Adds the source VLAN SNMP ifIndex to the information in the full-flow mask.

Release 12.2SRB

Cisco IOS Release 12.2SRB use the following flow masks:

- destination-source-interface—A more-specific flow mask. Adds the source VLAN SNMP ifIndex to the information in the destination-source flow mask.
- full-interface—The most-specific flow mask. Adds the source VLAN SNMP ifIndex to the information in the full-flow mask.

Other flow masks are handled as follows in order to accommodate per-interface mode for IPv4 traffic:

- Source-only, destination, and destination-source flow masks are treated as destination-source-interface.
- Full flow masks are treated as full-interface.

NDE Versions

NDE on the PFC supports NDE versions 5, 7, and 9 for the statistics captured on the PFC.
For information about NetFlow version 9, see the publication at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123_1/nfv9expf.htm

The following tables describe the supported NDE fields:

- Table 56-1—Version 5 header format
- Table 56-2—Version 7 header format

- Table 56-3—Version 5 flow record format
- Table 56-4—Version 7 flow record format

Table 56-1 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 56-2 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
20–23	reserved	Unused (zero) bytes

Note Some fields in the flow records might not have values, depending on the current flow mask. Unsupported fields contain a zero (0).

Table 56-3 NDE Version 5 Flow Record Format

			Flow masks: **• X=Populated** **• A=Additional field**					
Bytes	**Content**	**Description**	**Source**	**Destination**	**Destination Source**	**Destination Source Interface**	**Full**	**Full Interface**
0–3	srcaddr	Source IP address	**X**	**0**	**X**	**X**	**X**	**X**
4–7	dstaddr	Destination IP address	**0**	**X**	**X**	**X**	**X**	**X**
8–11	nexthop	Next hop router’s IP address[1]	**0**	**A**[2]	**A**	**A**	**A**	**A**

Table 56-3 NDE Version 5 Flow Record Format (continued)

Bytes	Content	Description	Flow masks: • X=Populated • A=Additional field					
			Source	Destination	Destination Source	Destination Source Interface	Full	Full Interface
12–13	input	Ingress interface SNMP ifIndex	0	0	0	X	0	X
14–15	output	Egress interface SNMP ifIndex[3]	0	A[2]	A	A	A	A
16–19	dPkts	Packets in the flow	X	X	X	X	X	X
20–23	dOctets	Octets (bytes) in the flow	X	X	X	X	X	X
24–27	first	SysUptime at start of the flow (milliseconds)	X	X	X	X	X	X
28–31	last	SysUptime at the time the last packet of the flow was received (milliseconds)	X	X	X	X	X	X
32–33	srcport	Layer 4 source port number or equivalent	0	0	0	0	X[4]	X[4]
34–35	dstport	Layer 4 destination port number or equivalent	0	0	0	0	X	X
36	pad1	Unused (zero) byte	0	0	0	0	0	0
37	tcp_flags	Cumulative OR of TCP flags[5]	0	0	0	0	0	0
38	prot	Layer 4 protocol (for example, 6=TCP, 17=UDP)	0	0	0	0	X	X
39	—	—	—	—	—	—	—	—
40–41	src_as	Autonomous system number of the source, either origin or peer	X	0	X	X	X	X
42–43	dst_as	Autonomous system number of the destination, either origin or peer	0	X	X	X	X	X
44–45	src_mask	Source address prefix mask bits	X	0	X	X	X	X
46–47	dst_mask	Destination address prefix mask bits	0	X	X	X	X	X
48	pad2	Pad 2	0	0	0	0	0	0

[1] Always zero when PBR, WCCP, or SLB is configured.

[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.

[3] Always zero when policy-based routing is configured.

[4] With PFC3CXL, PFC3C, PFC3BXL, or PFC3B, for ICMP traffic, contains the ICMP code and type values.

[5] Always zero for hardware-switched flows.

Table 56-4 NDE Version 7 Flow Record Format

Bytes	Content	Description	**Flow masks:** **• X=Populated** **• A=Additional field**					
			Source	**Destination**	**Destination Source**	**Destination Source Interface**	**Full**	**Full Interface**
0–3	srcaddr	Source IP address	X	0	X	X	X	X
4–7	dstaddr	Destination IP address	0	X	X	X	X	X
8–11	nexthop	Next hop router's IP address[1]	0	A[2]	A	A	A	A
12–13	input	Ingress interface SNMP ifIndex	0	0	0	X	0	X
14–15	output	Egress interface SNMP ifIndex[3]	0	A[2]	A	A	A	A
16–19	dPkts	Packets in the flow	X	X	X	X	X	X
20–23	dOctets	Octets (bytes) in the flow	X	X	X	X	X	X
24–27	First	SysUptime at start of the flow (milliseconds)	X	X	X	X	X	X
28–31	Last	SysUptime at the time the last packet of the flow was received (milliseconds)	X	X	X	X	X	X
32–33	srcport	Layer 4 source port number or equivalent	0	0	0	0	X[4]	X[4]
34–35	dstport	Layer 4 destination port number or equivalent	0	0	0	0	X	X
36	flags	Flow mask in use	X	X	X	X	X	X
37	tcp_flags	Cumulative OR of TCP flags[5]	0	0	0	0	0	0
38	prot	Layer 4 protocol (for example, 6=TCP, 17=UDP)	0	0	0	0	X	X
39	—	—	—	—	—	—	—	—
40–41	src_as	Autonomous system number of the source, either origin or peer	X	0	X	X	X	X
42–43	dst_as	Autonomous system number of the destination, either origin or peer	0	X	X	X	X	X
44	src_mask	Source address prefix mask bits	X	0	X	X	X	X
45	dst_mask	Destination address prefix mask bits	0	X	X	X	X	X
46–47	pad2	Pad 2	0	0	0	0	0	0
48–51	MLS RP	IP address of MLS router	0	X	X	X	X	X

[1] Always zero when PBR, WCCP, or SLB is configured.

[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.

[3] Always zero when policy-based routing is configured.

[4] With PFC3CXL, PFC3C, PFC3BXL, or PFC3B, for ICMP traffic, contains the ICMP code and type values.

[5] Always zero for hardware-switched flows.

MLS Cache Entries

NetFlow captures traffic statistics in the NetFlow cache on the PFC.

NetFlow maintains traffic statistics for each active flow in the NetFlow cache and increments the statistics when packets within each flow are switched. Periodically, NDE exports summarized traffic statistics for all expired flows, which the external data collector receives and processes.

Exported NetFlow data contains statistics for the flow entries in the NetFlow cache that have expired since the last export. Flow entries in the NetFlow cache expire and are flushed from the NetFlow cache when one of the following conditions occurs:

- The entry ages out.
- The entry is cleared by the user.
- An interface goes down.
- Route flaps occur.

To ensure periodic reporting of continuously active flows, entries for continuously active flows expire at the end of the interval configured with the **mls aging long** command (default 1920 seconds [32 minutes]).

NDE packets go to the external data collector either when the number of recently expired flows reaches a predetermined maximum or after:

- 30 seconds for version 5 export.
- 10 seconds for version 9 export.

By default, all expired flows are exported unless they are filtered. If you configure a filter, NDE only exports expired and purged flows that match the filter criteria. NDE flow filters are stored in NVRAM and are not cleared when NDE is disabled. See the "Configuring NDE Flow Filters" section on page 56-27 for NDE filter configuration procedures.

NetFlow Sampling

NetFlow sampling is used when you want to report statistics for a subset of the traffic flowing through your network. The Netflow statistics can be exported to an external collector for further analysis.

There are two types of NetFlow sampling; NetFlow traffic sampling and NetFlow flow sampling. The configuration steps for configuring MSFC-based NetFlow traffic sampling for traffic switched in the software path and PFC/DFC-based NetFlow flow sampling for traffic switched in the hardware path on a Cisco 7600 series router use different commands because they are mutually independent features.

The following sections provide additional information on the two types of NetFlow sampling supported by Cisco 7600 series routers:

- NetFlow Traffic Sampling, page 56-7
- NetFlow Flow Sampling, page 56-8

NetFlow Traffic Sampling

NetFlow traffic sampling provides NetFlow data for a subset of traffic forwarded by a Cisco router by analyzing only one randomly selected packet out of n sequential packets (n is a user-configurable parameter) from the traffic that is processed by the router. NetFlow traffic sampling is used on platforms that perform software-based NetFlow accounting, such as Cisco 7200 series routers and Cisco 7600 series MSFCs, to reduce the CPU overhead of running NetFlow by reducing the number of packets that are analyzed (sampled) by NetFlow. The reduction in the number of packets sampled by NetFlow on platforms that perform software based NetFlow accounting also reduces the number of packets that need to be exported to an external collector. Reducing the number of packets that need to be exported to an external collector by reducing the number of packets that are analyzed is useful when the volume of exported traffic created by analyzing every packet will overwhelm the collector, or result in an over-subscription of an outbound interface.

NetFlow traffic sampling and export for software-based NetFlow accounting behaves in the following manner:

- The flows are populated with statistics from a subset of the traffic that is seen by the router.
- The flows are expired.
- The statistics are exported.

On Cisco 7600 series routers, NetFlow traffic sampling is supported only on the MSFC for software switched packets. For more information on configuring NetFlow traffic sampling, see the *Cisco IOS NetFlow Configuration Guide*.

NetFlow Flow Sampling

NetFlow flow sampling does not limit the number of packets that are analyzed by NetFlow. NetFlow flow sampling is used to select a subset of the flows processed by the router for export. Therefore, NetFlow flow sampling is not a solution to reduce oversubscribed CPUs or oversubscribed hardware NetFlow table usage. NetFlow flow sampling can help reduce CPU usage by reducing the amount of data that is exported. Using NetFlow flow sampling to reduce the number of packets that need to be exported to an external collector by reporting statistics on only a subset of the flows is useful when the volume of exported traffic created by reporting statistics for all of the flows will overwhelm the collector, or result in an over-subscription of an outbound interface.

NetFlow flow sampling is available on Cisco 7600 series routers for hardware-based NetFlow accounting on the PFCs and DFCs installed in the router.

NetFlow flow sampling and export for hardware-based NetFlow accounting behaves in the following manner:

- Packets arrive at the switch and flows are created/updated to reflect the traffic seen.
- The flows are expired.
- The flows are sampled to select a subset of flows for exporting.
- The statistics for the subset of flows that have been selected by the NetFlow flow sampler are exported.

Note When NetFlow flow sampling is enabled, aging schemes such as fast, normal, long aging are disabled.

You can configure NetFlow flow sampling to use time-based sampling or packet-based sampling. With either the full-interface or destination-source-interface flow masks, you can enable or disable NetFlow Flow Sampling on each Layer 3 interface.

Packet-based NetFlow Flow Sampling

Packet-based NetFlow flow sampling uses a sampling-rate in packets and an interval in milliseconds to select a subset (sample) of flows from the total number of flows processed by the router. The values for the sampling-rate are: 64, 128, 256, 512, 1024, 2048, 4096, 8192. The interval is a user-configurable value in the range 8000-16000 milliseconds. The default for the interval is 16000 milliseconds. The

interval value replaces the aging schemes such as fast, normal, long aging for expiring flows from the cache. The command syntax for configuring packet-based NetFlow flow sampling is:
mls sampling packet-based *rate* [*interval*].

Packet-based NetFlow flow sampling uses one of these two methods to select flows for sampling and export:

- **The number of packets in the expired flow exceeds the sampling rate**: If in a interval of X - where X is a value in the range of 8000-16000 (inclusive), a flow has a greater number of packets than the value configured for the sampling-rate, the flow is sampled (selected) and then exported.
- **The number of packets in the expired flow is less than the sampling rate**: If in a interval of X - where X is a value in the range of 8000-16000 (inclusive), a flow has a smaller number of packets than the value configured for the sampling-rate, the packet count for the flow is added to one of eight buckets based on the number of packets in the flow. The eight bucket sizes are 1/8th increments of the sampling rate. The packet count for a flow that contains a quantity of packets that is 0–1/8th of the sampling rate is assigned to the first bucket. The packet count for a flow that contains a quantity of packets that is 1/8th–2/8th of the sampling rate is assigned to the second bucket. And so on. When adding the packet count for a flow to a bucket causes the counter for the bucket to exceed the sampling rate, the last flow for which the counters were added to the bucket is sampled and exported. The bucket counter is changed to 0 and the process of increasing the bucket counter is started over. This method ensures that some flows for which the packet count never exceeds the sampling rate are selected for sampling and export.

Time-based Netflow Flow Sampling

Time-based Netflow flow sampling samples flows created in the first sampling time (in milliseconds) of the export interval time (in milliseconds). Each of the sampling rates that you can configure with the **mls sampling time-based** *rate* command has fixed values for the sampling time and export interval used by time-based NetFlow flow sampling. For example:

- If you configure a sampling rate of 64, NefFlow flow sampling selects flows created within the first 64 milliseconds (sampling time) of every 4096 millisecond export interval.
- If you configure a sampling rate of 2048, NefFlow flow sampling selects flows created within the first 4 milliseconds (sampling time) of every 8192 millisecond export interval.

Table 56-5 lists the sampling rates and export intervals for time-based NetFlow flow sampling.

Table 56-5 Time-Based Sampling Rates, Sampling Times, and Export Intervals

Sampling Rate (Configurable)	Sampling Time in Milliseconds (Not Configurable)	Export Interval Milliseconds (Not Configurable)
1 in 64	64	4096
1 in 128	32	4096
1 in 256	16	4096
1 in 512	8	4096
1 in 1024	4	4096
1 in 2048	4	8192
1 in 4096	4	16384
1 in 8192	4	32768

NetFlow Aggregation

For information about NetFlow aggregation support on the PFC and DFCs, see the "NetFlow Aggregation" section of the document at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnfov.htm

Note

- In Release 12.2SRB, you must enable NetFlow on individual interfaces in order to enable the hardware flow cache to be populated. When enabled, the cache is populated with flows only from those interfaces where NetFlow is enabled.
- In Release 12.2SRA, configuring an aggregation scheme allows the hardware flow cache to be populated. The cache is globally populated with information for all Layer 3 interfaces.
- Configuring NetFlow aggregation for the MSFC also configures it for the PFC and DFCs. (See Configuring NetFlow Aggregation for Flows on the MSFC, page 56-23, for a pointer to configuration instructions).
- NetFlow aggregation uses NDE version 8.

Per-Interface NetFlow and NDE

In Cisco IOS Release 12.2SRB and later releases, the per-interface NetFlow and NDE feature allows you to enable NetFlow on individual interfaces in order to gather and export statistics for IPv4 traffic flows on those interfaces. Previously, when you enabled NetFlow, statistics were gathered for all of the interfaces on the router (global mode).

If you upgrade to Release 12.2SRB (per-interface mode) from an earlier release (global mode), you must issue the **ip flow ingress** command on individual interfaces to activate NetFlow. The upgrade process automatically converts existing global-mode flowmasks into the corresponding per-interface type (source, destination, and destination-source become destination-source-interface, and full becomes full-interface).

If you downgrade from Release 12.2SRB to an earlier release, NetFlow resumes global-mode operation (gathering statistics for all router interfaces) and the 12.2SRB flowmasks remain in effect.

The per-interface NetFlow feature improves NetFlow table utilization and performance as follows:

- Provides more room in the NetFlow table for flows that are of interest. With per-interface NetFlow, table entries are created only for those interfaces where NetFlow is enabled. This reduces the number of unwanted entries in the table, leaving more room for those flows that you are interested in. Previously, entries were created for all router interfaces.

  Creating table entries only for interfaces where NetFlow is enabled improves performance because:

  - The NetFlow table is shared by all flow-based features (NetFlow, QoS, multicast, and so on).
  - If the NetFlow table gets too full, NetFlow shortcuts might not being installed, which can result in flow statistics (and accounting information) being lost.
- Helps to ensure that the export of NDE records to the Netflow Data Collector (NFC) at a high rate of speed does not overwhelm the NFC and cause important accounting data to be lost. Since statistics are gathered and exported for specific interfaces only, the number of NDE records sent to the NFC is more manageable.
- Helps to ensure that there is less unintentional conflict between NDE and other features.

The following sections provide information about per-interface NetFlow and NDE and some additional NetFlow and NDE related features that are being introduced in Release 12.2SRB:

- Per-Interface NetFlow and NDE Usage Guidelines and Limitations, page 56-11
- Configuring Per-Interface NetFlow and NDE, page 56-12
- Verifying Per-Interface NetFlow and NDE, page 56-12
- NetFlow v9 for IPv6, page 56-13
- NDE on VRF Interfaces, page 56-13

Per-Interface NetFlow and NDE Usage Guidelines and Limitations

Consider the following usage guidelines and limitations when you configure per-interface NetFlow and NDE on the Cisco 7600 router:

- Supported in Cisco IOS Release 12.2SRB and later releases.
- Supported on RSP720, Sup720, and Sup32.
- Supported for IPv4 unicast and multicast traffic on Layer 3 interfaces.
  For IPv6 flows, NetFlow and NDE operate in global mode, not per-interface mode.
- When you enable NetFlow and NDE for Layer 2 (bridged) flows, the features are also automatically enabled for Layer 3 (routed) flows on the interface. To disable NetFlow and NDE for the interface, you must disable the feature for both the Layer 2 and Layer 3 flows. Use the **no ip flow ingress layer2-switched** command to disable L2 flows and **no ip flow ingress** to disable L3 flows.
- You can configure per-interface NetFlow and QoS micro-policing on an interface. However, do not configure different flow mask types on an interface. Only a single flow mask type should be configured for per-interface NetFlow and microflow policy.
- Beginning in Release 12.2SRB, the router supports both NDE flow mask and QoS flow mask; however, you cannot configure both types of flow masks on the same interface.
- When NDE and multicast non-RPF are both enabled, NDE has the potential to lose statistics. This potential loss occurs because NetFlow and NDE are enabled globally for multicast flows, which means that the NetFlow table could overflow.
- When you use the **platform ip features sequential** command on an interface , you must configure the interface-full flowmask feature. This enables the NDE to export the correct statistics, and avoids double accounting.
- The following limitations apply to flow masks in per-interface mode:
  - You cannot configure different flow mask types for individual interfaces. Only a single flow mask type is supported for all interfaces configured for per-interface NetFlow or NDE.
  - The same flow mask is used for both routed (L3) and bridged (L2) NetFlow entries for NDE.
  - All source and destination flow masks are treated as destination-source-interface and both of the full masks are treated as full-interface. See the "Flow Masks" section on page 56-3 for a description of flow mask types.
- All of guidelines and limitations in the "NetFlow and NDE Configuration Guidelines and Restrictions" section on page 56-14 apply.

Configuring Per-Interface NetFlow and NDE

Following is a summary of the steps you must perform to configure per-interface NetFlow and NDE on Cisco 7600 routers. Detailed procedures for each step are provided in the sections later in this chapter.

1. If you plan to export NetFlow statistics, globally enable NDE on the router by issuing the following commands:

```
configure terminal
ip flow-export destination
ip flow-export version
mls nde sender version
```

2. Enable NetFlow on individual interfaces by issuing the following commands:

```
configure terminal
interface
  ip flow ingress
```

3. (Optional) To configure NetFlow sampling, do the following:
   a. Enable sampled NetFlow globally on the router (refer to Configuring NetFlow Flow Sampling, page 56-17).
   b. Enable sampled NetFlow on individual interfaces (**mls netflow sampling**).
   c. Apply the config on the interface (**ip flow ingress**)
4. Verify the NDE configuration to ensure that it does not conflict with other features such as QoS or multicast. Use the **show ip interface** command to verify the configuration (see the "Verifying Per-Interface NetFlow and NDE" section on page 56-12).

Verifying Per-Interface NetFlow and NDE

To verify whether per-interface NetFlow and NDE are properly configured, use the **show ip interface** command (as shown here). In the command output, fields showing NetFlow and NDE configuration information are shown in boldface.

```
Router# show ip interface gig2/9
GigabitEthernet2/9 is up, line protocol is up
  Internet address is 10.0.0.1/8
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.2 224.0.0.6
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP CEF switching turbo vector
  IP Null turbo vector
```

```
Associated unicast routing topologies:
  Topology "base", operation state is UP
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Probe proxy name replies are disabled
Policy routing is disabled
Network address translation is disabled
BGP Policy Mapping is disabled
Input features: Ingress-NetFlow
Output features: Post-Ingress-NetFlow, HW Shortcut Installation
Post encapsulation features: HW Shortcut Installation
Sampled Netflow is disabled
IP Routed Flow creation is enabled in netflow table
IP Bridged Flow creation is disabled in netflow table
WCCP Redirect outbound is disabled
WCCP Redirect inbound is disabled
WCCP Redirect exclude is disabled
IP multicast multilayer switching is disabled
```

NetFlow v9 for IPv6

Cisco IOS Release 12.2SRB introduces support for NetFlow version 9 for IPv6. For information about how to configure this feature on the Cisco 7600 router, see its feature module description in the new feature documentation for Release 12.2SRB at the following URL:

http://www.cisco.com/en/US/products/ps6922/products_feature_guides_list.html

NDE on VRF Interfaces

Cisco IOS Release 12.2SRB introduces support for NDE on VRF interfaces. This new feature enables the Cisco 7600 router to capture and export NetFlow statistics for IPv4 packets in an MPLS Virtual Private Network (VPN). In this scenario, the router is functioning as provider edge (PE) router at the edge of an MPLS network.

For additional information about NDE on VRF interfaces, see its feature module description in the new feature documentation for Release 12.2SRB at the following URL:

http://www.cisco.com/en/US/products/ps6922/products_feature_guides_list.html

Default NetFlow and NDE Configuration

Table 56-6 shows the default NetFlow and NDE configuration.

Table 56-6   Default NetFlow and NDE Configuration

Feature	Default Value
NetFlow	Disabled. 12.2SRB—Per-interface mode for IPv4 unicast (global mode for all else). 12.2SRA—Global mode.
NDE	Disabled.
NDE on VRF interfaces	12.2SRB—Disabled. 12.2SRA—Not available.
NetFlow and NDE of ingress bridged IP traffic	Disabled.
NDE source addresses	None.
NDE data collector address and UDP port	None.
NDE filters	None.
NetFlow Mask	None.
NetFlow Sampling	Disabled.
NetFlow Aggregation	Disabled.
Populating additional NDE fields	Enabled.

NetFlow and NDE Configuration Guidelines and Restrictions

When configuring NetFlow and NDE, follow these guidelines and restrictions:

NetFlow and NDE support IP multicast traffic only with NetFlow version 9. With other NetFlow versions, you can display NetFlow statistics for IP multicast traffic with the **show mls ip multicast** command.

- Cisco 7600 routers do not support Netflow in egress direction for unicast ip packets.
- All PFCs (except the PFC3A) support NetFlow and NDE for bridged IP traffic.
- NDE does not support Internetwork Packet Exchange (IPX) traffic.
- The Policy Feature Card 3 (PFC3) does not use the NetFlow table for Layer 3 switching in hardware.
- If the NetFlow table utilization exceeds these recommended utilization levels, there is an increased probability that there will be insufficient room to store statistics:

PFC	Recommended NetFlow Table Utilization	Total NetFlow Table Capacity
PFC3CXL PFC3BXL	235,520 (230K) entries	262,144 entries
PFC3C PFC3B	117,760 (115K) entries	131,072 entries
PFC3A	65,536 (64K) entries	131,072 entries

- No statistics are available for flows that are switched when the NetFlow table is full.
- The Cisco 7600 series router uses the Netflow table to maintain information about flow-based features. Normally, the Feature Manager creates a Netflow table entry for a flow-based feature only on the line card where the flow ingresses. However, because TCP intercept is a global feature, the router creates an entry for each TCP intercept flow on each of the installed PFCs and DFCs, not just the ingress PFC or DFC. This means that the PFC or DFC where the TCP intercept flow ingresses will have a non-zero packet count, but the other PFC and the DFCs will have a count of zero packets for the flow. [CSCek47971]
- The following IPv4 Netflow and NDE options are not available for IPv6 flows: [CSCek55571]
  - Aggregation support (**ip flow-aggregation cache** command)
  - Export of Layer 2 switched IPv6 flows
  - Netflow and NDE sampling
  - NDE filter support

Multicast NDE Configuration Guidelines

Observe the following guidelines when you configure multicast NDE on the Cisco 7600:

- In Release 12.2SRB and later releases, multicast NDE and QoS microflow policing cannot both be configured on the same interface. However, the features can be configured on different interfaces.
- To configure multicast NDE, issue both the **ip flow ingress** and **ip multicast netflow ingress** commands. Note that the **ip multicast netflow ingress** command is enabled by default.

Release 12.2SRB and Later Releases

Beginning in Release 12.2SRB, for IPv4 flows, the router supports per-interface mode NetFlow and NDE only. For IPv6 flows, NetFlow and NDE continue to operate in global mode.

See the "Per-Interface NetFlow and NDE" section on page 56-10 for information about per-interface NetFlow and NDE and its usage guidelines and restrictions.

Configuring NetFlow and NDE

These sections describe how to configure NetFlow and NDE:

- Configuring NetFlow and NDE for Flows on the PFC, page 56-16
- Configuring NetFlow and NDE for Flows on the MSFC, page 56-22
- Enabling NetFlow and NDE for Ingress Bridged IP Traffic, page 56-24
- Displaying the NDE Address and Port Configuration, page 56-26

- Configuring NDE Flow Filters, page 56-27
- Displaying the NDE Configuration, page 56-29

Note

- You must enable NetFlow on the MSFC Layer 3 interfaces to support NDE on the PFC and NDE on the MSFC.
- You must enable NDE on the MSFC to support NDE on the PFC.
- When you configure NAT and NDE on an interface, the PFC sends all traffic in fragmented packets to the MSFC to be processed in software. (CSCdz51590)

Note NDE and NAT configuration on the same interface is not supported. NDE requires flows to age out periodicaly for it to export its statistics. NAT installs hardware shortcuts that do not age. Hence, NDE for NAT'd flows does not work correctly.

Configuring NetFlow and NDE for Flows on the PFC

These sections describe how to configure NetFlow and NDE for flows on the PFC:

- Configuring NetFlow for Flows on the PFC, page 56-16
- Enabling NDE, page 56-21

Configuring NetFlow for Flows on the PFC

These sections describe how to configure NetFlow statistics collection for flows on the PFC:

- Enabling NetFlow on the PFC (Release 12.2SRA), page 56-16
- Enabling Per-Interface NetFlow (Release 12.2SRB and Later), page 56-17
- Configuring NetFlow Flow Sampling, page 56-17
- Configuring NetFlow Aggregation for Flows on the PFC, page 56-19
- Setting the Minimum IP MLS Flow Mask (Release 12.2SRA Only), page 56-19
- Configuring the MLS Aging Time, page 56-20

Enabling NetFlow on the PFC (Release 12.2SRA)

To enable NetFlow statistics collection for flows on the PFC in Release 12.2SRA, perform this task. For information about enabling NetFlow in Release 12.2SRB and later releases, see the following section.

Command	Purpose
`Router(config)# mls netflow`	Enables NetFlow on the PFC.
`Router(config)# no mls netflow`	Disables NetFlow on the PFC.

This example shows how to enable NetFlow statistics collection:

```
Router(config)# mls netflow
```

Enabling Per-Interface NetFlow (Release 12.2SRB and Later)

To enable NetFlow statistics collection for flows on the PFC in Release 12.2SRB and later releases, perform this task. See the "Per-Interface NetFlow and NDE" section on page 56-10 for information about how the router operates in NetFlow and NDE per-interface mode. For detailed information about command syntax, see the command reference documents listed at the beginning of this chapter.

Command	Purpose
`Router(config)# mls flow ip`	Configures the flow mask to use for NetFlow entries.
`Router(config)# interface interface`	Selects the interface to enable NetFlow on.
`Router(config-if)# [no] ip flow ingress`	Enables NetFlow on a Layer 3 interface. Issue the command on each interface where you want to enable the feature. Use the **no** form of the command to disable NetFlow and NDE on the interface.
`Router(config-if)# exit`	Exits interface configuration mode.

`Router(config)# mls nde sender`
`Router(config)# ip flow-export destination {hostname | ip-address} udp-port`
	(Optional) Enables NDE. Issue these commands if you plan to export NetFlow statistics.
	Specifies an external host (name or IP address) to send NetFlow statistics to and the port

`Router(config)# mls nde sender`	(Optional) Enables NDE. Use this command if you plan to export NetFlow statistics.
`Router(config)# ip flow-export destination {hostname | ip-address} udp-port`	(Optional) Specifies the host name or IP address of the external host to export NetFlow statistics to and specifies the port to send the statistics to.
`Router(config)# show ip interface interface`	Displays the configuration of the specified interface. Examine the configuration to ensure that the NDE configuration does not conflict with other features such as QoS or multicast (see "Verifying Per-Interface NetFlow and NDE").

Configuring NetFlow Flow Sampling

These sections describe how to configure sampled NetFlow on the PFC:

- Configuring NetFlow Flow Sampling Globally (Release 12.2SRB and Release 12.2SRA), page 56-18
- Configuring Per-Interface Mode NetFlow Flow Sampling (Release 12.2SRB), page 56-18
- Configuring NetFlow Flow Sampling on a Layer 3 Interface (Release 12.2SRA), page 56-18

Note NDE on the MSFC does not support NetFlow Flow Sampling.

Configuring NetFlow Flow Sampling Globally (Release 12.2SRB and Release 12.2SRA)

To configure sampled NetFlow globally in Release 12.2SRB and Release 12.2SRA, perform this task:

	Command	Purpose
Step 1	`Router(config)# mls sampling {time-based rate | packet-based rate [interval]}`	Enables sampled NetFlow and configures the rate. For packet-based sampling, optionally configures the export interval.
	`Router(config)# no mls sampling`	Clears the sampled NetFlow configuration.
Step 2	`Router(config)# end`	Exits configuration mode.

When you configure sampled NetFlow globally, note the following information:

- The valid values for *rate* are 64, 128, 256, 512, 1024, 2048, 4096, and 8192.
- The valid values for the packet-based export *interval* are from 8,000 through 16,000.
- To export any data in Release 12.2SRA, you must also configure sampled NetFlow on a Layer 3 interface.

See the "NetFlow Sampling" section on page 56-7 for more information.

Configuring Per-Interface Mode NetFlow Flow Sampling (Release 12.2SRB)

In Release 12.2SRB and later releases, you must enable sampled NetFlow globally and on individual interfaces (as shown in the following example).

In the example, the **mls sampling** command enables sampled NetFlow globally and the **mls netflow sampling** command enables sampled NetFlow on the interface (in this example, Fast Ethernet port 5/12).

```
Router# configure terminal
Router(config)# mls sampling packet-based 64
Router(config)# interface fastethernet 5/12
Router(config-if)# mls netflow sampling
Router(config)# end
Router#
```

Configuring NetFlow Flow Sampling on a Layer 3 Interface (Release 12.2SRA)

In Release 12.2SRA, with the full-interface or destination-source-interface flow masks, you can enable or disable sampled NetFlow on individual Layer 3 interfaces. With all other flow masks, sampled NetFlow is enabled or disabled globally.

To configure sampled NetFlow on a Layer 3 interface in Release 12.2SRA, make sure that sampled NetFlow is enabled globally and perform this task:

	Command	Purpose
Step 1	`Router(config)# interface {vlan vlan_ID | type slot/port}`	Specifies the Layer 3 interface to configure. **Note** The Layer 3 interface must be configured with an IP address.
Step 2	`Router(config-if)# mls netflow sampling`	Enables sampled NetFlow on the Layer 3 interface.
	`Router(config-if)# no mls netflow sampling`	Disables sampled NetFlow on the Layer 3 interface.
Step 3	`Router(config)# end`	Exits configuration mode.

This example shows how to enable sampled NetFlow on Fast Ethernet port 5/12:

```
Router# configure terminal
Router(config)# interface fastethernet 5/12
Router(config-if)# mls netflow sampling
Router(config)# end
Router#
```

Configuring NetFlow Aggregation for Flows on the PFC

NetFlow aggregation is configured automatically for flows on the PFC and DFCs when you configure NetFlow aggregation for the MSFC (see the "Configuring NetFlow Aggregation for Flows on the MSFC" section on page 56-23 for a pointer to configuration instructions).

To display NetFlow aggregation cache information for the PFC or DFCs, perform this task:

Command	Purpose
Router # **show ip cache flow aggregation** {**as** | **destination-prefix** | **prefix** | **protocol-port** | **source-prefix**) **module** *slot_num*	Displays the NetFlow aggregation cache information.
Router # **show mls netflow aggregation flowmask**	Displays the NetFlow aggregation flow mask information. This command is applicable only to Release 12.2SRA; the command is not applicable in Release 12.2SRB.

Note The PFC and DFCs do not support NetFlow ToS-based router aggregation.

This example shows how to display the NetFlow aggregation cache information:

```
Router# show ip cache flow aggregation destination-prefix module 1
IPFLOW_DST_PREFIX_AGGREGATION records and statistics for module :1
IP Flow Switching Cache, 278544 bytes
2 active, 4094 inactive, 6 added
236 ager polls, 0 flow alloc failures
Active flows timeout in 30 minutes
Inactive flows timeout in 15 seconds
Dst If Dst Prefix Msk AS Flows Pkts B/Pk Active
Gi7/9 9.1.0.0 /16 0 3003 12M 64 1699.8
Gi7/10 11.1.0.0 /16 0 3000 9873K 64 1699.8
Router#
```

This example displays the NetFlow aggregation flow mask information (Release 12.2SRA only):

```
Router# show mls netflow aggregation flowmask
 Current flowmask set for netflow aggregation : Vlan Full Flow
 Netflow aggregations configured/enabled :
        AS Aggregation
        PROTOCOL-PORT Aggregation
        SOURCE-PREFIX Aggregation
        DESTINATION-PREFIX Aggregation
Router#
```

Setting the Minimum IP MLS Flow Mask (Release 12.2SRA Only)

You can set the minimum specificity of the flow mask for the NetFlow cache on the PFC (see the "Flow Masks" section on page 56-3). The actual flow mask that is used will have at least the specificity configured by the **mls flow ip** command.

Note The task does not apply to Release 12.2SRB, which supports only the interface-destination-source and interface-full flow masks.

To set the minimum IP flow mask, perform this task:

Command	Purpose
`Router(config)# mls flow ip {source | destination | destination-source | interface-destination-source | full | interface-full}`	Sets the minimum IP flow mask for the protocol.
`Router(config)# no mls flow ip`	Reverts to the default IP flow mask (null).

This example shows how to set the minimum IP flow mask:

```
Router(config)# mls flow ip destination
```

To display the IP flow mask configuration, perform this task:

Command	Purpose
`Router# show mls netflow flowmask`	Displays the flow mask configuration.

This example shows how to display the MLS flow mask configuration:

```
Router# show mls netflow flowmask
current ip flowmask for unicast: destination address
Router#
```

Configuring the MLS Aging Time

The MLS aging time (default 300 seconds) applies to all NetFlow cache entries. You can configure the normal aging time in the range of 32 to 4092 seconds. Flows can age as much as 4 seconds sooner or later than the configured interval. On average, flows age within 2 seconds of the configured value.

Other events might cause MLS entries to be purged, such as routing changes or a change in link state.

Note If the number of MLS entries exceeds the recommended utilization (see the "NetFlow and NDE Configuration Guidelines and Restrictions" section on page 56-14), only adjacency statistics might be available for some flows.

To keep the NetFlow cache size below the recommended utilization, enable the following parameters when using the **mls aging** command:

- **normal**—Configures the wait before aging out and deleting entries that are not covered by fast or long aging.
- **fast aging**—Configures an efficient process to age out entries created for flows that only switch a few packets, and then are never used again. The **fast aging** parameter uses the **time** keyword value to check if at least the **threshold** keyword value of packets have been switched for each flow. If a flow has not switched the threshold number of packets during the time interval, then the entry is aged out.

- **long**—Configures the aging time for deleting entries that are always in use. Long aging is used to prevent counter wraparound, which can cause inaccurate statistics.

A typical cache entry that is removed is the entry for flows to and from a Domain Name Server (DNS) or TFTP server. This entry might not be used again after it is created. The PFC saves space in the NetFlow cache for other data when it detects and ages out these entries.

If you need to enable MLS fast aging time, initially set the value to 128 seconds. If the size of the NetFlow cache continues to grow over the recommended utilization, decrease the setting until the cache size stays below the recommended utilization. If the cache continues to grow over the recommended utilization, decrease the normal MLS aging time.

To configure an MLS aging time, perform this task:

Command	Purpose
Router(config)# **mls aging** {**fast** [**threshold** {*1-128*} | **time** {*1-128*}] | **long** *64-1920* | **normal** *32-4092*}	Configures an MLS aging time for a NetFlow cache entry.
Router(config)# **no mls aging fast**	Disables fast aging.
Router(config)# **no mls aging** {**long** | **normal**}	Reverts to the default MLS aging time.

This example displays how to configure an MLS aging time:

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# mls aging fast threshold 64 time 30
```

To display the MLS aging-time configuration, perform this task:

Command	Purpose
Router# **show mls netflow aging**	Displays the MLS aging-time configuration.

This example shows how to display the MLS aging-time configuration:

```
Router# show mls netflow aging
enable timeout packet threshold
------ ------- ----------------
normal aging true 300 N/A
fast aging true 32 100
long aging true 900 N/A
```

Enabling NDE

For both Release 12.2SRA and Release 12.2 SRB, perform this task to globally enable NDE:

Command	Purpose
Router(config)# **mls nde sender** [**version** {**5** | **7**}]	Enables NDE for flows on the PFC and (optionally) configures the NDE version. Specify an NDE version that matches the NetFlow collector that the data is being exported to.
Router(config)# **ip flow-export destination** {*hostname* | *ip-address*} *udp-port*	Identifies the

Command	Purpose
Router(config)# **no mls nde sender**	Disables NDE for flows on the PFC.
Router(config)# **no mls nde sender version**	Reverts to the default (version 7).

Note

- NDE for the PFC uses the source interface configured for the MSFC (see the "Configuring the MSFC NDE Source Layer 3 Interface" section on page 56-23).
- NetFlow version 9 is supported—See this publication:

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123_1/nfv9expf.htm

This example shows how to globally enable NDE for flows on the PFC:

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

This example shows how to globally enable NDE for the PFC and configure NDE version 5:

```
Router(config)# mls nde sender version 5
```

Configuring NetFlow and NDE for Flows on the MSFC

This section supplements the NetFlow procedures at this URL:

http://www.cisco.com/en/US/docs/ios/12_2/switch/configuration/guide/xcfnfc.html

These sections describe how to configure NDE on the MSFC:

- Enabling NetFlow for Flows on the MSFC, page 56-22
- Configuring NetFlow Aggregation for Flows on the MSFC, page 56-23
- Configuring the MSFC NDE Source Layer 3 Interface, page 56-23
- Configuring the NDE Destination, page 56-24

Enabling NetFlow for Flows on the MSFC

In Release 12.2SRB and later releases, NDE is automatically enabled on an interface when you enable NetFlow on the interface (**ip flow ingress**). However, for NDE to work, you must globally enable it and specify a destination to export the statistics to (**mls nde sender** and **ip flow-export destination**).

In Release 12.2SRA, enable NetFlow on the MSFC by performing this task for each Layer 3 interface where you want to enable NDE.

	Command	Purpose
Step 1	Router(config)# **interface** {**vlan** *vlan_ID*} | {*type slot/port*} | {**port-channel** *port_channel_number*}	Selects a Layer 3 interface to configure.
Step 2	Router(config-if)# **ip flow ingress** Router(config-if)# **ip route-cache flow**	Enables NetFlow.

Note If Netflow is enabled on the port channel, then theflow entries are created per port-channel interface. NetFlow entries are not created for each port channel member link and the NetFlow from member links will be part of the port-channel NetFlow.

Configuring NetFlow Aggregation for Flows on the MSFC

To configure NetFlow aggregation for flows on the MSFC, use the procedures in the section "Configuring an Aggregation Cache" at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnfc.htm#wp1001058

Note

- Configuring NetFlow aggregation for the MSFC automatically configures it for the PFC and DFCs.
- In Release 12.2SRB, you must enable NetFlow on individual interfaces in order to enable the hardware flow cache to be populated. When enabled, the cache is populated with flows only from those interfaces where NetFlow is enabled.
- In Release 12.2SRA, configuring an aggregation scheme allows the hardware flow cache to be populated. The cache is globally populated with information for all L3 interfaces.

To configure NetFlow ToS-based router aggregation for the MSFC, use the procedures at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s15/dtnfltos.htm

Note The PFC and DFCs do not support NetFlow ToS-based router aggregation.

Configuring the MSFC NDE Source Layer 3 Interface

To configure the Layer 3 interface used as the source of the NDE packets containing statistics from the MSFC, perform this task:

Command	Purpose
`Router(config)# ip flow-export source {{vlan vlan_ID} | {type slot/port} | {port-channel number} | {loopback number}}`	Configures the interface used as the source of the NDE packets containing statistics from the MSFC.
`Router(config)# no ip flow-export source`	Clears the NDE source interface configuration.

When configuring the MSFC NDE source Layer 3 interface, note the following information:

- You must select an interface configured with an IP address.
- You can use a loopback interface.

This example shows how to configure a loopback interface as the NDE flow source:

```
Router(config)# ip flow-export source loopback 0
Router(config)#
```

Configuring the NDE Destination

To configure the destination IP address and UDP port to receive the NDE statistics, perform this task:

Command	Purpose
Router(config)# **ip flow-export destination** *ip_address udp_port_number*	Configures the NDE destination IP address and UDP port.
Router(config)# **no ip flow-export destination** *ip_address udp_port_number*	Clears the NDE destination configuration.

Note Netflow Multiple Export Destinations—To configure redundant NDE data streams, which improves the probability of receiving complete NetFlow data, you can enter the **ip flow-export destination** command twice and configure a different destination IP address in each command. This hardware supports the Netflow Multiple Export Destinations feature:

- PFC3

This example shows how to configure the NDE flow destination IP address and UDP port:

```
Router(config)# ip flow-export destination 172.20.52.37 200
```

Note The destination address and UDP port number are saved in NVRAM and are preserved if NDE is disabled and reenabled or if the router is power cycled. If you are using the NetFlow FlowCollector application for data collection, verify that the UDP port number you configure is the same port number shown in the FlowCollector's /opt/csconfc/config/nfconfig.file file.

Enabling NetFlow and NDE for Ingress Bridged IP Traffic

All PFCs (except the PFC3A) support NetFlow and NDE for ingress bridged IP traffic. The following sections describe how to enable NetFlow and NDE for ingress bridged IP traffic:

- Enabling NetFlow for Ingress Bridged IP Traffic in VLANs, page 56-25
- Enabling NDE for Ingress Bridged IP Traffic in VLANs, page 56-25

Note

- When you enable NetFlow for ingress bridged IP traffic, the statistics are available to the Sampled Netflow feature (see the "NetFlow Sampling" section on page 56-7).
- For each VLAN where you want to enable NetFlow and NDE for bridged IP traffic, you must create a corresponding VLAN interface, assign an IP address to it, and issue the **no shutdown** command to bring the interface up.
- When you enable NetFlow for bridged IP traffic on a VLAN, export of the bridged traffic is enabled by default as long as NDE is globally enabled.

Enabling NetFlow for Ingress Bridged IP Traffic in VLANs

To enable NetFlow for ingress bridged IP traffic in VLANs, perform this task:

Command	Purpose
`Router(config)# ip flow ingress layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]`	Enables NetFlow for ingress bridged IP traffic in the specified VLANs. **Note** NetFlow for ingress bridged IP traffic in a VLAN requires that NetFlow on the PFC be enabled with the **mls netflow** command.
`Router(config)# no ip flow ingress layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]`	Disables NetFlow for ingress bridged IP traffic in the specified VLANs.

This example shows how to enable NetFlow for ingress bridged IP traffic in VLAN 200:

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# ip flow ingress layer2-switched vlan 200
```

Enabling NDE for Ingress Bridged IP Traffic in VLANs

To enable NDE for ingress bridged IP traffic in VLANs, perform this task:

Command	Purpose
`Router(config)# ip flow export layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]`	Enables NDE for ingress bridged IP traffic in the specified VLANs (enabled by default when you enter the **ip flow ingress layer2-switched vlan** command). **Note** NDE for ingress bridged IP traffic in a VLAN requires that NDE on the PFC be enabled with the **mls nde sender** command.
`Router(config)# no ip flow export layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]`	Disables NDE for ingress bridged IP traffic in the specified VLANs.

This example shows how to enable NDE for ingress bridged IP traffic in VLAN 200:

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)# ip flow export layer2-switched vlan 200
```

Displaying the NDE Address and Port Configuration

To display the NDE address and port configuration, perform these tasks:

Command	Purpose
Router# **show mls nde**	Displays the NDE export flow IP address and UDP port configuration.
Router# **show ip flow export**	Displays the NDE export flow IP address, UDP port, and the NDE source interface configuration.

This example shows how to display the NDE export flow source IP address and UDP port configuration:

```
Router# show mls nde
Netflow Data Export enabled
 Exporting flows to  10.34.12.245 (9999)
 Exporting flows from 10.6.58.7 (55425)
 Version: 7
 Include Filter not configured
 Exclude Filter is:
   source:  ip address 11.1.1.0, mask 255.255.255.0
 Total Netflow Data Export Packets are:
    49 packets, 0 no packets, 247 records
 Total Netflow Data Export Send Errors:
        IPWRITE_NO_FIB = 0
        IPWRITE_ADJ_FAILED = 0
        IPWRITE_PROCESS = 0
        IPWRITE_ENQUEUE_FAILED = 0
        IPWRITE_IPC_FAILED = 0
        IPWRITE_OUTPUT_FAILED = 0
        IPWRITE_MTU_FAILED = 0
        IPWRITE_ENCAPFIX_FAILED = 0
 Netflow Aggregation Enabled
   source-prefix aggregation export is disabled
   destination-prefix aggregation exporting flows to 10.34.12.245 (9999)
10.34.12.246 (9909)
      exported 84 packets, 94 records
   prefix aggregation export is disabled
Router#
```

This example shows how to display the NDE export flow IP address, UDP port, and the NDE source interface configuration:

```
Router# show ip flow export
Flow export is enabled
  Exporting flows to 172.20.52.37 (200)
  Exporting using source interface FastEthernet5/8
  Version 1 flow records
  0 flows exported in 0 udp datagrams
  0 flows failed due to lack of export packet
  0 export packets were sent up to process level
  0 export packets were dropped due to no fib
  0 export packets were dropped due to adjacency issues
Router#
```

Configuring NDE Flow Filters

These sections describe NDE flow filters:

- NDE Flow Filter Overview, page 56-27
- Configuring a Port Flow Filter, page 56-27
- Configuring a Host and Port Filter, page 56-27
- Configuring a Host Flow Filter, page 56-28
- Configuring a Protocol Flow Filter, page 56-28

NDE Flow Filter Overview

By default, all expired flows are exported until you configure a filter. After you configure a filter, only expired and purged flows matching the specified filter criteria are exported. Filter values are stored in NVRAM and are not cleared when NDE is disabled.

To display the configuration of the NDE flow filters you configure, use the **show mls nde** command described in the "Displaying the NDE Configuration" section on page 56-29.

Configuring a Port Flow Filter

To configure a destination or source port flow filter, perform this task:

Command	Purpose
Router(config)# **mls nde flow** {**exclude** | **include**} {**dest-port** *number* | **src-port** *number*}	Configures a port flow filter for an NDE flow.
Router(config)# **no mls nde flow** {**exclude** | **include**}	Clears the port flow filter configuration.

This example shows how to configure a port flow filter so that only expired flows to destination port 23 are exported (assuming the flow mask is set to full):

```
Router(config)# mls nde flow include dest-port 23
Router(config)#
```

Configuring a Host and Port Filter

To configure a host and TCP/UDP port flow filter, perform this task:

Command	Purpose
Router(config)# **mls nde flow** {**exclude** | **include**} {**destination** *ip_address mask* | **source** *ip_address mask* {**dest-port** *number* | **src-port** *number*}}	Configures a host and port flow filter for an NDE flow.
Router(config)# **no mls nde flow** {**exclude** | **include**}	Clears the port flow filter configuration.

This example shows how to configure a source host and destination TCP/UDP port flow filter so that only expired flows from host 171.69.194.140 to destination port 23 are exported (assuming the flow mask is set to ip-flow):

```
Router(config)# mls nde flow include source 171.69.194.140 255.255.255.255 dest-port 23
```

Configuring a Host Flow Filter

To configure a destination or source host flow filter, perform this task:

Command	Purpose
Router(config)# **mls nde flow** {**exclude** | **include**} {**destination** *ip_address mask* | **source** *ip_address mask* | **protocol** {**tcp** {**dest-port** *number* | **src-port** *number*} | **udp** {**dest-port** *number* | **src-port** *number*}}	Configures a host flow filter for an NDE flow.
Router(config)# **no mls nde flow** {**exclude** | **include**}	Clears port filter configuration.

This example shows how to configure a host flow filter to export only flows to destination to host 172.20.52.37:

```
Router(config)# mls nde flow include destination 172.20.52.37 255.255.255.255
Router(config)#
```

Configuring a Protocol Flow Filter

To configure a protocol flow filter, perform this task:

Command	Purpose
Router(config)# **mls nde flow** {**exclude** | **include**} **protocol** {**tcp** {**dest-port** *number* | **src-port** *number*} | **udp** {**dest-port** *number* | **src-port** *number*}}	Configures a protocol flow filter for an NDE flow.
Router(config)# **no mls nde flow** {**exclude** | **include**}	Clears port filter configuration.

This example shows how to configure a TCP protocol flow filter so that only expired flows from destination port 35 are exported:

```
Router(config)# mls nde flow include protocol tcp dest-port 35
Router(config)#
```

To display the status of the NDE flow filters, use the **show mls nde** command described in the "Displaying the NDE Configuration" section on page 56-29.

Usage Guidelines to Configure Protocol Flow Filter

Follow these restrictions and usage guidelines to configure NetFlow Data Export Filter:

- Only one filter is supported to include or exclude flow export. The flow export configuration is based on source IP, destination IP, source Port, destination port and protocol.
- If you separately configure each filter parameter, the final filter consists of all the configured filter values as shown in the next example:

```
Router(config)#mls nde flow include src-port 100
Router#sh run | I mls nde flow
mls nde flow include protocol tcp src-port 100
Router(config)#mls nde flow include dest-port 200
Router#sh run | I mls nde flow
mls nde flow include protocol tcp src-port 100 dest-port 200
Router#
```

- If you reconfigure a filter with a new value, the old value is overwritten as shown in the next example:

```
Router(config)#mls nde flow include dest-port 200
Router#sh run | I mls nde flow
mls nde flow include dest-port 200
Router(config)#mls nde flow include dest-port 500
Router#sh run | I mls nde flow
mls nde flow include dest-port 500
```

Displaying the NDE Configuration

To display the NDE configuration, perform this task:

Command	Purpose
`Router# show mls nde`	Displays the NDE configuration.

This example shows how to display the NDE configuration:

```
Router# show mls nde
 Netflow Data Export enabled
 Exporting flows to  10.34.12.245 (9988)  10.34.12.245 (9999)
 Exporting flows from 10.6.58.7 (57673)
 Version: 7
 Include Filter not configured
 Exclude Filter not configured
 Total Netflow Data Export Packets are:
    508 packets, 0 no packets, 3985 records
 Total Netflow Data Export Send Errors:
        IPWRITE_NO_FIB = 0
        IPWRITE_ADJ_FAILED = 0
        IPWRITE_PROCESS = 0
        IPWRITE_ENQUEUE_FAILED = 0
        IPWRITE_IPC_FAILED = 0
        IPWRITE_OUTPUT_FAILED = 0
        IPWRITE_MTU_FAILED = 0
        IPWRITE_ENCAPFIX_FAILED = 0
 Netflow Aggregation Enabled
Router#
```

NetFlow Support on GRE Tunnels

This section describes implementation of NetFlow accounting for IPv4 unicast flows over GRE tunnels on 7600 platform.

GRE is a tunneling protocol developed by Cisco and it is capable of encapsulating a wide variety of protocol packet types within IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork This feature is supported on 15.3(3)S4. IPv4 unicast flows can be exported in any NDE format.

Configuration

“ip flow ingress” needs to be configured on both the physical interface (f1/1) and tunnel interface (t0) for the encapsulated flows to be accounted. The following are the examples of ip flow ingress configuration.

Configuring NetFlow Accounting on a GRE IP Tunnel Example - Encapsulation Node

```
mls flow ip interface - full
mls sampling packet - based 64 8000
interface Tunnel1
ip address 12.0.0.1 255.255.255.252 12.0.0.2 255.255.255.252
ip flow ingress
mls netflow sampling
tunnel source 172.1.0.1
tunnel destination 172.1.0. 172.1.0.2
```

Configuring NetFlow Accounting on a GRE IP Tunnel Example - Decapsulation Node

```
mls flow ip interface - full
mls sampling packet - based 64 8000
interface Tunnel1
ip address 12.0.0.2 255.255.255.252
ip flow ingress
mls netflow sampling
tunnel source 172.1.0.1
tunnel destination 172.1.0.2
```

Accounting of Flows

Unicast IPv4 packets encapsulated in IPv4 unicast GRE tunnels - this case needs to be considered for NetFlow accounting of packets in and out of a GRE tunnel.

The following figures explain the different terminologies used in the architecture

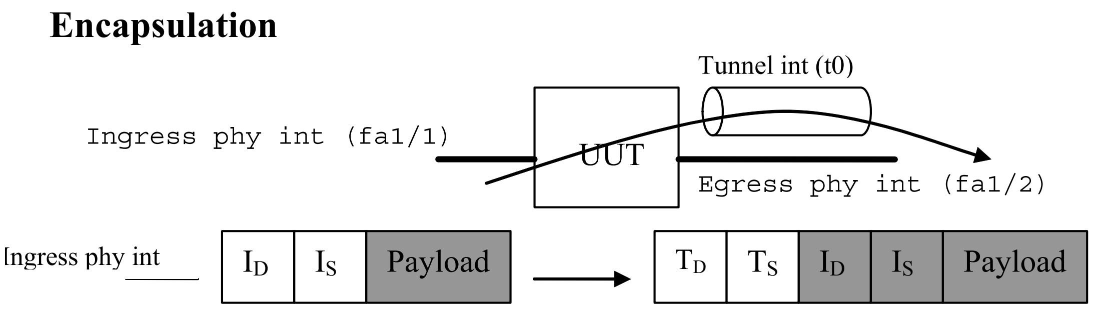

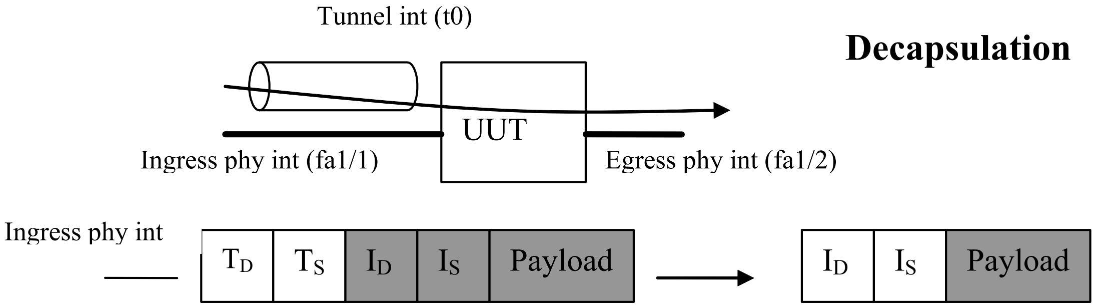

The following table represents the configuration on different interfaces and the flows that need to be created.

	Ingress Tunnel Interface	Ingress Physical Interface	Egress Physical Interface	Egress Tunnel Interface	Flows
Unicast Over GRE (Encap)	Ip flow ingress		No configuration	IP flow ingress	Flow (1) [Id, Is, fa1/1, t0] Flow (2) [Td, Ts, t0, fa1/2]
Unicast Over GRE (Decap)	Ip flow ingress	IP flow ingress	No configuration		Flow (1) [Td, Ts, fa1/1, t0] Flow (2) [Id, Is, t0, fa1/2]

The following display output shows that NetFlow accounting is operational because these are statistics for the hardware-switched NetFlow flows.

```
CE1#sh mls netflow ip module 5
Displaying Netflow entries in module 5
DstIP           SrcIP           Prot:SrcPort:DstPort  Src i/f
:AdjPtr
-----------------------------------------------------------------------------
Pkts         Bytes         Age   LastSeen  Attributes
------------------------------------------------------
12.0.0.1        12.0.0.2        icmp:0      :0        Tu1              :0x0
0            0             1     08:35:20   L3 - Dynamic
40.0.0.1        50.0.0.1        icmp:8      :0        Tu1              :0x0
0            0             1     08:35:20   L3 - Dynamic
```

Impact on Memory and Performance

The flows that get encapsulated or decapsulated on the router will now create two or more flows. This will have an impact on the hardware NetFlow table as more number of flows will be created. These flows need to be exported as well. Flows from the hardware table are converted to different format internally and then exported.

As the number of flows gets doubled, the required memory also gets doubled to convert the flows to different format internally.

Limitations

Though it supports 7600 with both sup720 and RSP720, it does not support mGRE and IPv6 packets in GRE. This design only supports hardware switched flows. The packets which are software switched are processed by the software path. This feature is supported only on ES+ line cards.

The following are the limitations of the updated design:

- NetFlow accounting on secondary VLAN of GRE tunnel is not supported
- In case secondary VLAN is present, accounting will happen only at the decapsulation side of the tunnel.
- This feature supports only P2P GRE tunnels
- MPLS aware NetFlow is not supported
- Number of flows depend on TCAM size
- Software - Feature is only applicable for GRE tunnel with unicast traffic
- Hardware - Only 7600 is supported