



Configuring NDE

This chapter describes how to configure NetFlow Data Export (NDE).



Note

For complete syntax and usage information for the commands used in this chapter, refer to these publications:

- The *Cisco IOS NetFlow Command Reference* at this URL:
http://www.cisco.com/en/US/docs/ios/netflow/command/reference/nf_book.html
- The Release 12.2 publications at this URL:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1838/tsd_products_support_series_home.html
- NetFlow version 9 is supported—See this document:
Cisco IOS NetFlow Configuration Guide.

This chapter contains the following sections:

- [Understanding NDE, page 51-1](#)
- [NDE Configuration Guidelines and Restrictions, page 51-10](#)
- [Configuring NDE, page 51-10](#)

Understanding NDE

These sections describe how NetFlow Data Export (NDE) works:

- [NDE Overview, page 51-1](#)
- [NDE on the MSFC, page 51-2](#)

NDE Overview

NetFlow collects traffic statistics by monitoring packets that flow through the router and storing the statistics in the NetFlow table. For more information about NetFlow, see [Chapter 50, “Configuring NetFlow”](#).

NetFlow Data Export (NDE) converts the NetFlow table statistics into records and exports the records to an external device, which is called a NetFlow collector.

In PFC3A mode, NDE exports statistics only for routed traffic. With a PFC3B or PFC3BXL, you can configure NDE to export statistics for both routed and bridged traffic. Netflow for bridged traffic requires Release 12.2(18)SXE or later.

You can export IP unicast statistics using NDE record format versions 5, 7 or 9. Use NDE version 8 record format for NetFlow aggregation, and version 9 record format for IP multicast. NetFlow version 9 export format is supported in Release 12.2(18)SXF and later.

Exporting a large volume of statistics can significantly impact SP and RP CPU utilization. You can control the volume of records exported by configuring NDE flow filters to include or exclude flows from the NDE export. When you configure a filter, NDE exports only the flows that match the filter criteria.

You can configure up to two external data collector addresses. A second data collector improves the probability of receiving complete NetFlow data by providing redundant data streams. This feature is available with the following releases and hardware:

- A PFC2 and Release 12.2(18)SXD and later releases
- A PFC3 and Release 12.2(18)SXE and later releases

NDE on the MSFC

NDE on the MSFC exports statistics for flows routed in software. The MSFC supports NetFlow aggregation, described in this document:

Cisco IOS NetFlow Configuration Guide.

The MSFC also supports NetFlow ToS-based router aggregation, described in this document:

Cisco IOS NetFlow Configuration Guide.

Release 12.2(18)SXF and later releases support NetFlow sampling on the MSFC, described in this document:

Cisco IOS NetFlow Configuration Guide.

Release 12.2(18)SXF and later releases support NetFlow version 9, described in this document:

Cisco IOS NetFlow Configuration Guide.

NetFlow version 9 record formats are described in this document:

Cisco IOS NetFlow Configuration Guide.

NDE on the PFC

NDE on the PFC exports statistics for flows routed or bridged in hardware. These sections describe NDE on the PFC in more detail:

- [NDE Flow Mask, page 51-3](#)
- [NDE Versions, page 51-3](#)
- [Exporting NetFlow Data, page 51-7](#)
- [NetFlow Sampling, page 51-7](#)

NDE Flow Mask

You can configure the minimum NetFlow flow mask for NDE. The NetFlow flow mask determines the granularity of the statistics gathered, which controls the volume of statistics for NDE to export.

For more details about flow masks, refer to [Chapter 50, “Configuring NetFlow”](#).

Additional NDE Fields

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
- BGP AS

These fields are populated by the software looking up the FIB table entry before sending out the NDE record to the collector. Therefore, these fields are blank when you use the **show** command to display the hardware NetFlow table.

NDE Versions

Release 12.2(18)SXF and later releases support NetFlow version 9, described at this URL:

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

NDE exports statistics for NetFlow aggregation flows using NDE version 8. The following document describes the version 8 header format:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgr/fswtch_c/swprt2/xcfnfov.htm#wp1001212

NDE exports IP unicast traffic using NDE versions 5, 7 and 9.

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

The following tables describe the supported fields for NDE versions 5 and 7:

- [Table 51-1](#)—Version 5 header format
- [Table 51-2](#)—Version 7 header format
- [Table 51-3](#)—Version 5 flow record format
- [Table 51-4](#)—Version 7 flow record format

NetFlow version 9 record formats are described at this URL:

http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00805e395a.html#wp1180857

Table 51-1 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 51-2 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 |

Table 51-3 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 51-12) | | | | | |
|-------|-----------|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|----------------|--------------------|------------------------------|----------------|----------------|
| | | | 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 | 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 | tos | IP type-of-service byte | X ⁶ | X ⁶ | X ⁶ | X ⁶ | X ⁶ | X ⁶ |
| 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. In PFC3BXL or PFC3B mode, for ICMP traffic, contains the ICMP code and type values.
5. Always zero for hardware-switched flows.
6. Populated in PFC3BXL or PFC3B mode.

Table 51-4 NDE Version 7 Flow Record Format

| Bytes | Content | Description | Flow masks: • X=Populated • A=Additional field (see the “Populating Additional NDE Fields” section on page 51-12) | | | | | |
|-------|-----------|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|----------------|--------------------|------------------------------|----------------|----------------|
| | | | 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 | tos | IP type-of-service byte | X ⁶ | X ⁶ | X ⁶ | X ⁶ | X ⁶ | X ⁶ |
| 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. In PFC3BXL or PFC3B mode, for ICMP traffic, contains the ICMP code and type values.
5. Always zero for hardware-switched flows.
6. Populated with Release 12.2(17b)SXA and later releases in PFC3BXL or PFC3B mode.

Exporting NetFlow Data

NetFlow maintains traffic statistics for each active flow in the NetFlow table 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 table that have expired since the last export. Flow entries in the NetFlow table expire and are flushed from the NetFlow table 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 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 51-16](#) 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 6500 series switch 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 6500 series switches:

- [NetFlow Traffic Sampling, page 51-7](#)
- [NetFlow Flow Sampling, page 51-8](#)

NetFlow Traffic Sampling

NetFlow traffic sampling provides NetFlow data for a subset of traffic forwarded by a Cisco router or switch 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 or switch. NetFlow traffic sampling is used on platforms that perform software-based NetFlow accounting, such as Cisco 7200 series routers and Cisco 6500 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 6500 series switches, 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 Catalyst 6500 series switches 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, NetFlow 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, NetFlow flow sampling selects flows created within the first 4 milliseconds (sampling time) of every 8192 millisecond export interval.

Table 51-5 lists the sampling rates and export intervals for time-based NetFlow flow sampling.

Table 51-5 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 |

Default NDE Configuration

Table 51-4 shows the default NDE configuration.

Table 51-6 Default NDE Configuration

| Feature | Default Value |
|-----------------------------------------|---------------|
| NDE | Disabled |
| NDE of ingress bridged IP traffic | Disabled |
| NDE source addresses | None |
| NDE data collector address and UDP port | None |
| NDE filters | None |
| Populating additional NDE fields | Enabled |

NDE Configuration Guidelines and Restrictions

When configuring NDE, follow these guidelines and restrictions:

- NDE supports IP multicast traffic only with [NetFlow version 9](#).
- NetFlow aggregation must use NDE version 8 or version 9.
- In PFC3B or PFC3BXL mode with Release 12.2(18)SXE and later releases, NDE supports bridged IP traffic. PFC3A mode does not support NDE for bridged IP traffic.
- NDE does not support Internetwork Packet Exchange (IPX) traffic or any other non-IP protocol.

Configuring NDE

These sections describe how to configure NDE:

- [Configuring NDE on the PFC, page 51-11](#)
- [Configuring NDE on the MSFC, page 51-13](#)
- [Enabling NDE for Ingress-Bridged IP Traffic, page 51-15](#)
- [Displaying the NDE Address and Port Configuration, page 51-15](#)
- [Configuring NDE Flow Filters, page 51-16](#)
- [Displaying the NDE Configuration, page 51-18](#)



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 fragmented packets to the MSFC to be processed in software. (CSCdz51590)

Configuring NDE on the PFC

These sections describe how to configure NDE on the PFC:

- [Enabling NDE From the PFC, page 51-11](#)
- [Populating Additional NDE Fields, page 51-12](#)
- [Configuring NetFlow Flow Sampling, page 51-12](#)

Enabling NDE From the PFC

To enable NDE from the PFC, perform this task:

| Command | Purpose |
|-----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Router(config)# mls nde sender [version {5 7}] | Enables NDE from the PFC using version 7 records or version 5 records. If you enter the mls nde sender command without using the version {5 7} keywords version 7 records are enabled by default. |
| Router(config)# ip flow-export version 9 | (Optional) Enables the use of version 9 records ¹ . If you want to enable the use of version 9 records for NDE, you must enter the mls nde sender command first. |
| Router(config)# no ip flow-export version 9 | Note Enabling the use of version 9 records overrides the use of either version 5 records or version 7 records. |
| Router(config)# no mls nde sender | Disables the use of version 9 records. Disables NDE from the PFC. |

1. The **ip flow-export version 9** command was integrated into 12.2(18)SXF.



Note

- NDE from the PFC uses the source interface configured for the MSFC (see the “[Configuring the MSFC NDE Source Layer 3 Interface](#)” section on page 51-13).
- Release 12.2(18)SXF and later releases support NetFlow version 9, described at this URL:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123_1/nfv9expf.htm

This example shows how to enable NDE from the PFC:

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

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

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

Populating Additional NDE Fields

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
- BGP AS

Not all of the additional fields are populated with all flow masks. See the [“NDE Versions” section on page 51-3](#) 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
```

Configuring NetFlow Flow Sampling

These sections describe how to configure NetFlow flow sampling on the PFC:

- [Configuring NetFlow Flow Sampling Globally, page 51-12](#)
- [Configuring NetFlow Flow Sampling on a Layer 3 Interface, page 51-13](#)

Configuring NetFlow Flow Sampling Globally

To configure NetFlow flow sampling globally, perform this task:

| | Command | Purpose |
|---------------|----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|
| Step 1 | Router(config)# mls sampling { time-based <i>rate</i> packet-based <i>rate</i> [<i>interval</i>]} | 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 NetFlow flow sampling 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.
- With a PFC3, to export any data, you must also configure sampled NetFlow on a Layer 3 interface.

Configuring NetFlow Flow Sampling on a Layer 3 Interface



Note

- With the full-interface or destination-source-interface flow masks, you can enable or disable NetFlow flow sampling on individual Layer 3 interfaces. With all other flow masks, NetFlow flow sampling is enabled or disabled globally.
- The Layer 3 interface must be configured with an IP address.

To configure NetFlow flow sampling on a Layer 3 interface, perform this task:

| | Command | Purpose |
|--------|------------------------------------------------------------------|----------------------------------------------------------|
| Step 1 | Router(config)# interface {vlan vlan_ID type slot/port} | Selects a Layer 3 interface to configure. |
| Step 2 | Router(config-if)# mls netflow sampling | Enables NetFlow flow sampling on the Layer 3 interface. |
| | Router(config-if)# no mls netflow sampling | Disables NetFlow flow sampling on the Layer 3 interface. |
| Step 3 | Router(config)# end | Exits configuration mode. |

This example shows how to enable NetFlow flow sampling 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 NDE on the MSFC

These sections describe how to configure NDE on the MSFC:

- [Configuring the MSFC NDE Source Layer 3 Interface, page 51-13](#)
- [Configuring the NDE Destination, page 51-14](#)
- [Configuring Netflow Flow Sampling, page 51-14](#)

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 <i>ip_address</i> <i>udp_port_number</i> | Configures the NDE destination IP address and UDP port. |
| Router(config)# no ip flow-export destination <i>ip_address</i> <i>udp_port_number</i> | 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 and these releases support the NetFlow Multiple Export Destinations feature:

- PFC3 and Release 12.2(18)SXE and later releases
- PFC2 and Release 12.2(18)SXD and later releases

Note that configuring two destinations increases the RP CPU utilization, as you are exporting the data records twice.

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.

Configuring Netflow Flow Sampling

In 12.2(18)SXF and later releases, the MSFC supports NetFlow sampling for software-routed traffic.

For additional information, see the following document:

Cisco IOS NetFlow Configuration Guide.

Enabling NDE for Ingress-Bridged IP Traffic

In PFC3B or PFC3BXL mode with Release 12.2(18)SXE and later releases, NDE supports ingress-bridged IP traffic. PFC3A mode does not support NDE for bridged IP traffic.



Note

To enable NetFlow for bridged IP traffic on a VLAN, you must create a corresponding VLAN interface, assign it an IP address, and enter the **no shutdown** command to bring up the interface. The exported bridged flows will have ingress and egress VLAN information and not the physical port information.

NDE is enabled by default when you enable NetFlow on the VLAN. To disable NDE for ingress-bridged IP traffic in VLANs, perform this task:

| Command | Purpose |
|-----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>Router(config)# ip flow export layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]</pre> | <p>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).</p> <p>Note NDE for ingress-bridged IP traffic in a VLAN requires that NDE on the PFC be enabled with the mls nde sender command.</p> |
| <pre>Router(config)# no ip flow export layer2-switched vlan vlan_ID[-vlan_ID] [, vlan_ID[-vlan_ID]]</pre> | <p>Disables NDE for ingress-bridged IP traffic in the specified VLANs.</p> |

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 |
|----------------------------------------|-------------------------------------------------------------------------------------------------------|
| <pre>Router# show mls nde</pre> | <p>Displays the NDE export flow IP address and UDP port configuration.</p> |
| <pre>Router# show ip flow export</pre> | <p>Displays the NDE export flow IP address, UDP port, and the NDE source interface configuration.</p> |

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 51-16](#)
- [Configuring a Port Flow Filter, page 51-17](#)
- [Configuring a Host and Port Filter, page 51-17](#)
- [Configuring a Host Flow Filter, page 51-17](#)
- [Configuring a Protocol Flow Filter, page 51-18](#)

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 51-18.

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 <i>number</i> src-port <i>number</i> } | 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 <i>ip_address mask</i> source <i>ip_address mask</i> { dest-port <i>number</i> src-port <i>number</i> }} | 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 <i>ip_address mask</i> source <i>ip_address mask</i> protocol { tcp { dest-port <i>number</i> src-port <i>number</i> } udp { dest-port <i>number</i> src-port <i>number</i> }} | 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 host 172.20.52.37:

```
Router(config)# mls nde flow include destination 172.20.52.37 255.255.255.225
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 <i>number</i> src-port <i>number</i> } udp { dest-port <i>number</i> src-port <i>number</i> }} | 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 51-18.

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#
```