



IPFIX

Table 1: Feature History Table

Feature Name	Release Information	Description
IPFIX Flow Record Enhancements for L2 and L3 traffic.	Release 7.4.1	This release introduces: <ul style="list-style-type: none">• Support for flow-based IPFIX protocol version 10(v10), for L2 interfaces. Only L3 interfaces were supported in previous releases.• A new record-type, MPLS-IPv4, to capture BGP next-hop information.

Internet Protocol Flow Information Export (IPFIX) is an IETF standard export protocol for sending Netflow packets. IPFIX is based on Netflow version 9.

The IPFIX feature formats Netflow data and transfers the Netflow information from an exporter to a collector using UDP as transport protocol.

Restrictions for IPFIX

These IPFIX features are not supported:

- Variable-length information element in the IPFIX template
- Stream Control Transmission Protocol (SCTP) as the transport protocol

Limitations for IPFIX


- You cannot modify an exporter version of an exporter map that is already applied to an interface. To modify the exporter version, first remove the exporter configuration applied on the interface, later modify the version and apply the configuration to the interface.
- An interface can have eight different monitor-maps but all the monitor maps should have the same version for the exporters. There can be different exporters for the 8 monitor maps but they all need to have the same exporter version either v9 or IPFIX.

- [Configuring IPFIX, on page 2](#)
- [IP Flow Information Export \(IPFIX\) 315, on page 9](#)

Configuring IPFIX

Consider SP-PE use case where SP (Service Provider) cloud is connected to the PE (Provider Edge) router through TenGigabit ethernet.

Figure 1: SP-PE Topology



Configuring NetFlow on PE router involves:

1. Configuring Exporter map with IPFIX as an exporter
2. Configuring Monitor map
3. Configuring Sampler map
4. Applying the Monitor map and Sampler map to an interface

Configuring Exporter map with IPFIX as the exporter version

```
flow exporter-map fem_ipfix
 destination 10.1.1.1
 source Loopback 0
 transport udp 1025
 exit
version ipfix
 template data timeout 600
 options sampler-table
 exit
```

Configuring Monitor map

```
flow monitor-map fmml
 record ipv4
 option filtered
 exporter fem_ipfix
 cache entries 10000
 cache timeout active 1800
```

```
cache timeout inactive 15
exit
```

Configuring Sampler map

```
sampler-map fsm1
  random 1 out-of 65535
exit
```

Applying the Monitor map to an interface

Now apply the monitor-map **fmm1** that is configured with an exporter version IPFIX and sampler-map **fsm1** to the 10GE 0/0/0/1 interface in the ingress direction:

```
configure
interface 10GE0/0/0/1
  flow ipv4 monitor fmm1 sampler fsm1 ingress
exit
```

Verification

Use the **show flow flow-exporter map** command to verify the exporter version configured is IPFIX:

```
RP/0/RSP0/CPU0:router# show flow exporter-map fem_ipfix
Flow Exporter Map : fem_ipfix
-----
Id                : 3
Packet-Length     : 1468
DestinationIpAddr : 10.1.1.1
VRFName           : default
SourceIfName      : Loopback1
SourceIpAddr      : 4.4.0.1
DSCP              : 40
TransportProtocol : UDP
TransportDestPort : 9001
```

Export Version: IPFIX

```
Common Template Timeout : 1800 seconds
Options Template Timeout : 1800 seconds
Data Template Timeout : 1800 seconds
Interface-Table Export Timeout : 0 seconds
Sampler-Table Export Timeout : 0 seconds
VRF-Table Export Timeout : 0 seconds
```

Exported packets in an IPFIX packet structure are in the form of template set or data set. The first data template is sent when the configuration is activated on the interface.

With constant stream, the flowset data does not change, so data is decoded. Data template is updated in the case of timeout on the template. To change the timeout options in the flow exporter, use the `template options timeout` command:

```
RP/0/RP0/CPU0:router(config)#flow exporter-map ipfix_exp1
RP/0/RP0/CPU0:router(config-fem)#version ipfix
RP/0/RP0/CPU0:router(config-fem-ver)#template options
RP/0/RP0/CPU0:TU-PE3(config-fem-ver)#template options timeout
RP/0/RP0/CPU0:TU-PE3(config-fem-ver)#template options timeout 30
```

```

RP/0/RP0/CPU0:router# show flow exporter-map ipfix_exp1
version ipfix

  template data timeout 30
!
dscp 40
transport udp 9001
source Loopback0
destination 10.127.59.86

```

IPFIX Enablement for SRv6 and Services over SRv6 Core

Table 2: Feature History Table

Feature Name	Release Information	Description
IPFIX Enablement for SRv6 and Services over SRv6 Core	Release 7.10.1	<p>During the transition from conventional IP/MPLS networks to SRv6-based networks, the necessity for monitoring SRv6 traffic flow becomes crucial. This feature enables IPFIX to effectively monitor SRv6 IP traffic flow from network devices.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> The srv6 keyword is introduced in the record ipv6 command. <p>The srv6 keyword is supported on fourth generation and later ASR 9000 Series High Density Ethernet line cards.</p>

Feature Name	Release Information	Description
Simultaneous L2 and L3 Flow Monitoring using IPFIX	Release 7.10.1	<p>This feature introduces support for simultaneous L2 and L3 flow monitoring. Now, you can configure IP Flow Information Export (IPFIX) to actively monitor and record end-to-end L2 and L3 flow information elements from network devices. Previously, only L2 or L3 flow could be monitored at a time.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> • The l2-l3 keyword is introduced in the record ipv4 command. • The l2-l3 keyword is introduced in the record ipv6 command. <p>YANG DATA models:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-UM-flow-cfg</code> (see Github, YANG Data Models Navigator) <p>The l2-l3 keyword is supported on fourth generation and later ASR 9000 Series High Density Ethernet line cards.</p>

During the transition from conventional IP/MPLS networks to SRv6-based networks, the requirement for information elements specific to SRv6 traffic flow arises. To address this requirement, we have introduced the **srv6** keyword within the **ipv6** command. Consequently, information related to SRv6 payload such as L2VPN and L3VPN services will also will be exported as part of IPFIX record.

Restriction and Limitation

1. IPFIX with multiple SRH is not supported in IOS XR software version 7.10.1
2. When the VLAN rewrite pop/translate option is enabled, the fourth and fifth generation of the Cisco ASR 9000 line cards do not support capturing of the VLAN information on an L2 interface.
3. SRv6 encapsulated L2VPN IPFIX records captured at the Decap PE node may show IE89 ForwardingStatus as "forwarded," but IE14 egressInterface will be 0.
4. When ASR 9000 is the endpoint of SR, Base Format 1 Segment Identifier (SID) is not supported and only the Micro-SID format for Layer 2 VPN services is supported.

Configuration

From Cisco IOS-XR Release 7.10.1, a new optional keyword, `srv6` is introduced for the `record ipv6` option. See the following example:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config-fem)# flow monitor-map MON
RP/0/RSP0/CPU0:router(config-fmm)# record ipv6 srv6
RP/0/RSP0/CPU0:router(config-fmm)# exporter EXP
RP/0/RSP0/CPU0:router(config-fmm)# cache timeout inactive 5
RP/0/RSP0/CPU0:router(config-fmm)# !
RP/0/RSP0/CPU0:router(config-fmm)# sampler-map SAMP
RP/0/RSP0/CPU0:router(config-fmm)# random 1 out-of 1000
RP/0/RSP0/CPU0:router(config-fmm)# !
RP/0/RSP0/CPU0:router(config-fmm)# interface GigabitEthernet0/1/0/0
RP/0/RSP0/CPU0:router(config-fmm)# ipv6 address 2002:1::1/64
RP/0/RSP0/CPU0:router(config-fmm)# flow ipv6 monitor M1 sampler SAMP ingres
```

This example shows how to display SRv6 monitor-map data for a specific flow:

```
RP/0/RSP0/CPU0:router# show flow monitor-map MON

Flow Monitor Map : MON
-----
Id:                1
RecordMapName:     srv6
ExportMapName:     EXP
CacheAgingMode:    Normal
CacheMaxEntries:   65535
CacheActiveTout:   1800 seconds
CacheInactiveTout: 5 seconds
CacheUpdateTout:   N/A
CacheRateLimit:    2000
HwCacheExists:     False
HwCacheInactTout:  50
```

From Cisco IOS-XR Release 7.10.1, a new optional keyword, `12-13` is introduced for the `record ipv4` and `record ipv6` option. See the following example:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config-fem)# flow monitor-map M-IPv4
RP/0/RSP0/CPU0:router(config-fmm)# record ipv4 12-13
RP/0/RSP0/CPU0:router(config-fmm)# exporter EXP-ipfix
RP/0/RSP0/CPU0:router(config-fmm)# !
RP/0/RSP0/CPU0:router(config-fmm)# flow monitor-map M-IPv6
RP/0/RSP0/CPU0:router(config-fmm)# record ipv6 12-13
RP/0/RSP0/CPU0:router(config-fmm)# exporter EXP-ipfix
RP/0/RSP0/CPU0:router(config-fmm)# !
RP/0/RSP0/CPU0:router(config-fmm)# sampler-map SAMP
RP/0/RSP0/CPU0:router(config-fmm)# random 1 out-of 1000
RP/0/RSP0/CPU0:router(config-fmm)# !
RP/0/RSP0/CPU0:router(config-fmm)# interface GigabitEthernet0/1/0/0
RP/0/RSP0/CPU0:router(config-fmm)# description CE-PE Interface
RP/0/RSP0/CPU0:router(config-fmm)# ipv4 address 1.1.1.1 255.255.255.0
RP/0/RSP0/CPU0:router(config-fmm)# ipv6 address 2001:DB8:c18:1::/64
RP/0/RSP0/CPU0:router(config-fmm)# flow ipv4 monitor M-IPv4 sampler SAMP ingres
RP/0/RSP0/CPU0:router(config-fmm)# flow ipv6 monitor M-IPv6 sampler SAMP ingress
RP/0/RSP0/CPU0:router(config-fmm)# !
RP/0/RSP0/CPU0:router
```

This example shows how to display IPv4 monitor-map data for a specific flow:

```
RP/0/RSP0/CPU0:router# show run flow monitor-map

flow monitor-map M-IPv4
  record ipv4 l2-13
  exporter EXP
!
flow monitor-map M-IPv6
  record ipv6 l2-13
  exporter EXP
!
```

This example shows how to display l2-l3 monitor-map data for IPv4 specific flow:

```
RP/0/RSP0/CPU0:router# show flow monitor-map M-IPv4

Flow Monitor Map : M-IPv4
-----
Id:                3
RecordMapName:     ipv4-l2-13
ExportMapName:     EXP
CacheAgingMode:    Normal
CacheMaxEntries:   65535
CacheActiveTout:  1800 seconds
CacheInactiveTout: 15 seconds
CacheUpdateTout:  N/A
CacheRateLimit:   2000
HwCacheExists:    False
HwCacheInactTout: 50
```

This example shows how to display l2-l3 monitor-map data for IPv6 specific flow:

```
RP/0/RSP0/CPU0:router# show flow monitor-map M-IPv6

Flow Monitor Map : M-IPv6
-----
Id:                4
RecordMapName:     ipv6-l2-13
ExportMapName:     EXP
CacheAgingMode:    Normal
CacheMaxEntries:   65535
CacheActiveTout:  1800 seconds
CacheInactiveTout: 15 seconds
CacheUpdateTout:  N/A
CacheRateLimit:   2000
HwCacheExists:    False
HwCacheInactTout: 50
```

This example shows the complete recorded data for SRv6 L2 services :

```
RP/0/RSP0/CPU0:router# show flow monitor M-IPv6 location 0/0/CPU0

Cache summary for Flow Monitor M1:
Cache size:                65535
Current entries:           3
Flows added:                4
Flows not added:           0
Ager Polls:                68143
```

```

- Active timeout                0
- Inactive timeout              1
- Immediate                     0
- TCP FIN flag                  0
- Emergency aged                0
- Counter wrap aged             0
- Total                         1
Periodic export:
- Counter wrap                  0
- TCP FIN flag                  0
Flows exported                   1

===== Record number: 1 =====
IPv6SrcAddr      : 2::2
IPv6DstAddr      : bbbb:bc00:88:e000::
BGPDstOrigAS     : 0
BGPSrcOrigAS     : 0
BGPNextHopV6    : fe80::232:17ff:fe7e:1ce1
IPv6TC           : 0
IPv6FlowLabel    : 50686
IPv6OptHdrs      : 0x0
IPV6Prot         : 143
L4SrcPort        : 0
L4DestPort       : 0
L4TCPFlags       : 0
IPV6DstPrfxLen  : 48
IPV6SrcPrfxLen  : 128
InputInterface   : Hu0/0/0/10
OutputInterface  : BE111.1
ForwardStatus    : Fwd
FirstSwitched   : 01 18:51:25:797
LastSwitched    : 01 18:51:25:797
ByteCount        : 61004304
PacketCount      : 113814
Dir              : Ing
SamplerID        : 1
InputVRFID       : default
OutputVRFID      : default
InnerIPV4SrcAddr : 0.0.0.0
InnerIPV4DstAddr : 0.0.0.0
InnerIPv6SrcAddr : ::
InnerIPv6DstAddr : ::
InnerL4SrcPort   : 0
InnerL4DestPort  : 0
SrcMacAddr       : 00:0c:29:0e:d8:32
DstMacAddr       : 00:0c:29:0e:d8:3c
EthType          : 2048
Dot1qPriority     : 0
Dot1qVlanId      : 2001
RecordType       : SRv6 L2 Service Record
SRHFlags         : 0x0
SRHTags          : 0x0
SRHSegmentsLeft  : 0
SRHNumSegments   : 0

```

This example shows the complete recorded data for IPv6 L2-L3 services :

```

RP/0/RSP0/CPU0:router# show flow monitor M-IPv6 location 0/0/CPU0

RP/0/RP0/CPU0:router# show flow monitor MON-MAP-v6 location 0/0/CPU0
Thu Apr 28 11:36:47.622 IST
...

```



```

===== Record number: 1 =====
IPv6SrcAddr      : 151:1::1
IPv6DstAddr      : ff02::1:ff00:2
BGPDstOrigAS    : 0
BGPSrcOrigAS    : 0
BGPNextHopV6    : ::
IPv6TC          : 224
IPv6FlowLabel    : 0
IPv6OptHdrs     : 0x0
IPv6Prot        : icmpv6
MinimumTTL      : 255
MaximumTTL      : 255
L4SrcPort       : 0
L4DestPort      : 135
L4TCPFlags      : 0
IPv6DstPrfxLen  : 0
IPv6SrcPrfxLen  : 0
InputInterface  : BE999.1
OutputInterface : 0
ForwardStatus   : FwdNoFrag
FirstSwitched   : 01 18:51:25:797
LastSwitched    : 01 18:51:25:797
ByteCount       : 104
PacketCount     : 1
Dir             : Ing
SamplerID       : 1
InputVRFID      : default
OutputVRFID     : default
SrcMacAddr      : 00:0c:29:0e:d8:32
DstMacAddr      : 00:0c:29:0e:d8:3c
EthType         : 2048
Dot1qPriority    : 0
Dot1qVlanId     : 100
CustVlanId      : 200

```

IP Flow Information Export (IPFIX) 315

Internet Protocol Flow Information Export (IPFIX) is an IETF standard export protocol (RFC 7011) for sending IP flow information. Cisco ASR 9000 Router supports IPFIX 315 format to export flow information. IPFIX 315 format facilitates sending ‘n’ octets frame information starting from ethernet header till transport header of the traffic flow over the network. IPFIX 315 supports sending variable size packet record with variable payload information such as IPv4, IPv6, MPLS, and Nested packets like OuterIP-GRE-InnerIP etc. The process includes sampling and exporting the traffic flow information. Along with the ethernet frame information, IPFIX 315 format exports information of incoming and outgoing interface of the sampled packet.

Use **hw-module profile netflow ipfix315 location** < *linecard location* > command to enable IPFIX 315.

The information of the packets flowing through a device is used for variety of purpose including network monitoring, capacity planning, traffic management, etc.

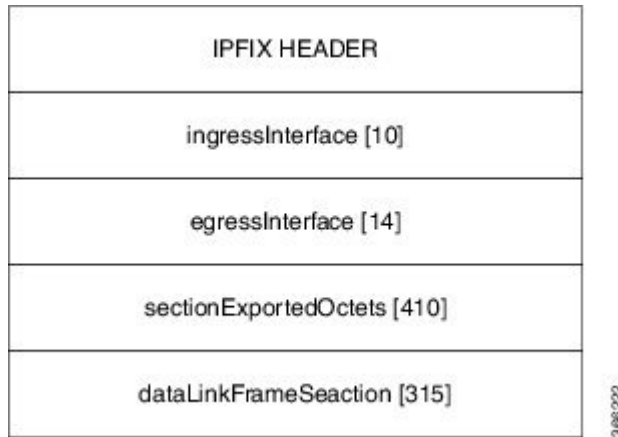
Sampling and Exporting Information

You must configure a sampling map to sample the traffic flow information. The sampler map specifies the rate at which packets (one out of n packets) are sampled.

The size of exported packet is until and including L4 header. If the L4 header is not found then the maximum of 160 bytes are exported.

The below figure *IPFIX 315 Export Packet Format* shows exported packet information.

Figure 2: IPFIX 315 Export Packet Format



A special cache type called Immediate Aging is used while exporting the packets. Immediate Aging ensures that the flows are exported as soon as they are added to the cache. Use the command **cache immediate** in flow monitor map configuration to enable Immediate Aging cache type.

IPFIX 315 Implementation Considerations

Here are few key points to consider before implementing IPFIX 315:

- You cannot enable the IPFIX 315 (using the `datalinkframesection` command) on an interface that has IPv4, IPv6 and MPLS flows already configured. Similarly, you cannot configure IPv4, IPv6 and MPLS flows if you have first enabled the IPFIX 315.
- Supported only in ingress direction.
- Supported on third and fourth generation of ASR 9000 line cards.
- Not supported on satellite interface.
- Supports only L3 routed packets.

Configuring IPFIX 315

Configuring IPFIX 315 involves:

1. Configuring Exporter map
2. Configuring Monitor map
3. Configuring Sampler map
4. Applying the Monitor map and Sampler map to an interface

Configuring Exporter map

```
flow exporter-map ipfix_exp
version ipfix
```

```

!
dscp 40
transport udp 9002
source Loopback1
destination 100.10.1.112
!

```



Note For **options** command and its configurations in Exporter Map, see [options](#).

Configuring Monitor map

```

flow monitor-map ipfix_mon
record datalinksectiondump
exporter ipfix_exp
cache immediate
cache entries 1000000
cache timeout rate-limit 1000000
!

```

Configuring Sampler map

```

sampler-map ipfix_sm
random 1 out-of 32000
!

```



Note The default cache size is 65535, hence you can configure sampling rate as 1 out of 65535 packets. However the recommended sampling rate is 1 out of 32000 packets.

Applying the Monitor map to an interface

```

interface HundredGigE 0/0/0/18
flow datalinkframesection monitor ipfix_mon sampler ipfix_sm ingress

```

Verification

Use the **show flow platform producer statistics location** command to display the statistics for datalinkframesection in the ingress direction:

```

RP/0/RP0/CPU0#show flow platform producer statistics location 02/CPU0
Wed Dec  6 02:49:04.411 EST
Netflow Platform Producer Counters:
IPv4 Ingress Packets:          3558922
IPv4 Egress Packets:           183
IPv6 Ingress Packets:          0
IPv6 Egress Packets:           0
MPLS Ingress Packets:         2176292132
MPLS Egress Packets:          96276772
Section Ingress Packets      2176292157
Drops (no space):              0
Drops (other):                 0
Unknown Ingress Packets:       0
Unknown Egress Packets:        0
Worker waiting:                 369792
SPP Packets:                    2119944979
Flow Packets:                   2276128009

```

```
Flow Packets per SPP Frame:          1
```

Use the **show flow monitor <monitor-map> cache location** command to check the flow monitor stats. In this example flow statistics for *ipfix_mon* monitor map are displayed:

```
RP/0/RP0/CPU0#show flow monitor ipfix_mon cache location 0/2/CPU0
```

```
Cache summary for Flow Monitor ipfix:
Cache size:                65535
Current entries:           0
Flows added:               2515
Flows not added:           0
Ager Polls:                252
- Active timeout           0
- Inactive timeout         0
- Immediate                2515
- TCP FIN flag             0
- Emergency aged           0
- Counter wrap aged        0
- Total                    2515
Periodic export:
- Counter wrap             0
- TCP FIN flag             0
Flows exported              2

Matching entries:          0
```

In the above sample output, cache immediate entries are 2515 and flows exported are 2.



Note The cache record statistics are not displayed for IPFIX 315.
