

Migrating from Traditional to Flexible NetFlow

Why Migrate to Flexible NetFlow?

If you are reading this document, you're either already convinced or curious about the potential advantages that Cisco® Flexible NetFlow will bring. For us at Cisco, this is a normal transition, and some new platforms and software releases will exclusively support Flexible NetFlow. As time goes on, we will never come back to Traditional NetFlow, so it's better to get prepared for the transition.

In this document we will call "Traditional NetFlow" everything that is not "Flexible NetFlow."

Some might simply think of using Flexible NetFlow the same way they have used Traditional NetFlow: that is, with the same flow record and by exporting with NetFlow Version 5 or NetFlow Version 9. Although this is possible and provides an interim solution to achieving a smooth migration that does not require any modification on existing collectors, we don't recommend going down that road, as it won't allow you to take advantage of the full capabilities of Flexible NetFlow.

Traditional NetFlow used a fixed seven tuple of IP information to identify a flow most of the time. A big advantage of the new Flexible NetFlow concept is that the user can define the flow. The benefits of Flexible NetFlow include:

- Flexible NetFlow will integrate with Network-Based Application Recognition (NBAR) to provide application visibility rather than just "flow" visibility. This positions Flexible NetFlow as a unique tool to differentiate and meter applications right from within the network.
- Because only interesting flows with selected key fields will be analyzed, Flexible NetFlow generally offers better performance, scalability, and aggregation of flow information.
- Enhanced flow infrastructure for security monitoring and distributed denial of service (DDoS) detection and identification.
- New information from packets to adapt flow information to a particular service or operation in the network. Flexible NetFlow users can customize the flow information available.
- Extensive use of Cisco's flexible and extensible NetFlow Version 9 export format.
- A comprehensive IP accounting feature that can be used to replace many accounting features, such as IP accounting, BGP Policy Accounting, and persistent caches.
- New high-end platforms such as Cisco Catalyst® 6000 with EARL8, Cisco Catalyst 4000 with K10, the next generation of Cisco Catalyst 3000, and so on will exclusively support Flexible NetFlow.

Traditional NetFlow allows you to understand what the network is doing and thus to optimize network design and reduce operational costs. With Flexible NetFlow the notion of "flow" goes beyond Layers 2/3/4. It gives you greater visibility and allows you to understand network behavior with more efficiency, with specific flow information tailored for various services used in the network.

What's Changing When Migrating to Flexible NetFlow?

We have tried to reduce the pain for you in making the transition from Traditional NetFlow to Flexible NetFlow; however, a few points might require more attention.

New Data Export Protocol

The export protocol of choice for Flexible NetFlow is the NetFlow Version 9 export protocol, but unfortunately and to date, NetFlow Version 5 has been a much more widely used protocol because of the legacy Cisco IOS® Software images that are still around that supported the NetFlow v5 export protocol only and worked very well. As mentioned in the previous section, Flexible NetFlow can also be configured to export some predefined flow records using the NetFlow Version 5 protocol format for backward compatibility.

As we make the transition to Flexible NetFlow's new model, one gains the ability to select key fields and non-key fields, to export many different fields (for example, packet fragments), to export Multiprotocol Label Switching (MPLS) labels or Border Gateway Protocol (BGP) next-hop fields, and so on. These fields cannot be transmitted over NetFlow Version 5 and can only be exported with a protocol that is as flexible as NetFlow Version 9, or later, IPFIX.

The main feature of the NetFlow Version 9 export format is that it is template-based. A template describes a NetFlow record format and the attributes of the fields (such as type and length) within the record. The router assigns each template an ID, which is communicated to the NetFlow Collection Engine along with the template description. The template ID is used for all further communication from the router to the NetFlow Collection Engine. (See Table 1.)

Table 1. Overview of Protocols per Version of NetFlow

NetFlow Metering Process	Information Elements	NetFlow Export Protocol	Transport Protocol
Traditional NetFlow	Traditional	NetFlow Versions 5 and 9	User Datagram Protocol (UDP)
	New (IPv6, Multicast)	NetFlow Version 9	UDP/Stream Control Transmission Protocol (SCTP)
Flexible NetFlow	Predefined Record for Traditional NetFlow	NetFlow Versions 5 and 9	UDP
	Other	NetFlow Version 9	UDP

Note: For the moment, Flexible NetFlow cannot export the flow information with SCTP (Reliable NetFlow Export) but only with UDP. If you really need SCTP export, you can still run Traditional NetFlow, exporting with SCTP at the same time Flexible NetFlow is running.

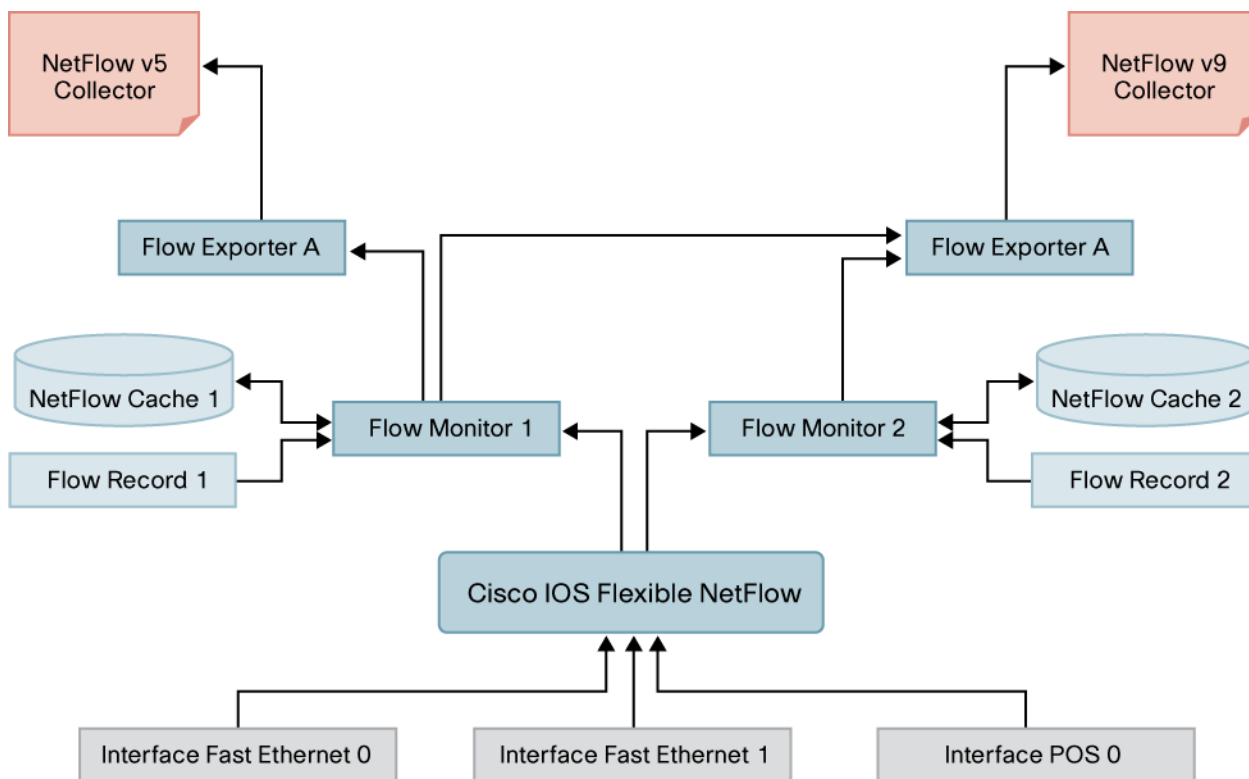
Introduction of a New Configuration CLI

Flexible NetFlow consists of components that can be used together in several variations to perform traffic analysis and data export, and the new command-line interface (CLI) configuration follows the same logic.

The user-defined flow records and the component structure of Flexible NetFlow make it easy for you to create various configurations for traffic analysis and data export on a networking device with a minimum number of configuration commands. Flow monitors can be defined according to the user's own requirements. Each flow monitor can have a unique combination of flow record, flow exporter, and cache type. If you change a parameter such as the destination IP address for a flow exporter, it is automatically changed for all the flow monitors that use the flow exporter. The same flow monitor can be used in conjunction with different flow samplers to sample the same type of network traffic at different rates on different interfaces. A single flow monitor can be attached to multiple interfaces, and multiple flow monitors can be attached to each interface.

Figure 1 shows how the information travels from the interfaces, through the processes, to collectors in Flexible NetFlow.

Figure 1. The Various Elements of Flexible NetFlow



The following sections provide more information on Flexible NetFlow components.

Using Records for Your Transition

In Flexible NetFlow a combination of key and non-key fields is called a flow record. Flow records are assigned to flow monitors to define the cache layout that is used to store the monitor's flow data. Flexible NetFlow includes several predefined records that can help you get started using Flexible NetFlow. To use Flexible NetFlow to its fullest potential, you should create your own customized records.

NetFlow Predefined Records

Flexible NetFlow includes several predefined records that you can use right away to start monitoring traffic in your network. The predefined records are available to help you quickly deploy Flexible NetFlow. You can choose from a list of already defined records that might meet the needs for network monitoring. Two of the predefined records (NetFlow original and NetFlow IPv4/IPv6 original output) emulate original NetFlow.

The predefined records help ensure backward compatibility with your existing NetFlow collector configurations for the data that is exported. Each of the predefined records has a unique combination of key and non-key fields that offer you the built-in ability to monitor various types of traffic in your network without customizing Flexible NetFlow on your router.

If you want to learn more about Flexible NetFlow predefined records, refer to the "Getting Started with Configuring Cisco IOS Flexible NetFlow" module or the "Configuring Cisco IOS Flexible NetFlow with Predefined Records" module.

User-Defined Records

With Flexible NetFlow you can define your own records for Flexible NetFlow flow monitor caches by specifying the key and non-key fields to customize the data collection to your specific requirements. When you define your own records for Flexible NetFlow flow monitor caches, they are referred to as user-defined records. The values of the key fields differentiate one flow from another and are taken from the first packet in the flow. The values in non-key fields are added to flows to provide additional information about the traffic in the flows. A change in the value of a non-key field does not create a new flow. In most cases the values for non-key fields are taken from only the first packet in the flow. However, exceptions are made for counters, flags, and min/max values. Flexible NetFlow lets you capture counter values such as the number of bytes and packets in a flow as non-key fields.

You can create user-defined records for applications such as quality of service (QoS) and bandwidth monitoring, application and end-user traffic profiling, and security monitoring for DoS attacks.

Example: Packet Section

Flexible NetFlow user-defined records provide the capability to monitor a contiguous section of a packet of a user-configurable size and use it in a flow record as a key or a non-key field along with other fields and attributes of the packet. The section might potentially include any Layer 3 data from the packet.

The packet section fields allow the user to monitor any packet fields that are not covered by the Flexible NetFlow predefined keys. The ability to analyze packet fields that are not collected with the predefined keys allows more detailed traffic monitoring, facilitates the investigation of DDoS attacks, and facilitates implementation of other security applications such as URL monitoring.

Flexible NetFlow adds a new Version 9 export format field type for the header and packet section types. Flexible NetFlow communicates to the NetFlow collector the configured section sizes in the corresponding Version 9 export template fields. The payload sections have a corresponding length field that can be used to collect the actual size of the collected section.

Flow Monitors

Flow monitors are the Flexible NetFlow components that are applied to interfaces to perform network traffic monitoring. Flow monitors consist of a user-defined or predefined record, an optional flow exporter, and a cache that is automatically created at the time the flow monitor is applied to the first interface. Flow data is collected from the network traffic and added to the flow monitor's cache during the monitoring process based on the key fields in the flow record.

Flexible NetFlow can be used to perform different types of analysis on the same traffic using an appropriate selection of key and non-key fields and, optionally, using different flow monitor cache types.

Three Flow Monitor Cache Types Instead of One

There are three types of flow monitor caches. You change the type of cache used by the flow monitor after you create the flow monitor. The three possible caches are:

Normal

The default cache type is "normal." In this mode, the entries in the cache are aged out according to the timeout active and timeout inactive settings. When a cache entry is aged out, it is removed from the cache and exported through any exporters configured.

Immediate

A cache of type "immediate" ages out every record as soon as it is created. As a result, every flow contains just one packet. The commands that display the cache contents will provide a history of the packets seen.

This mode is desirable when you expect only very small flows and you want a minimum amount of latency between seeing a packet and exporting a report.

Permanent

A cache of type "permanent" never ages out any flows. A permanent cache is useful when the number of flows you expect to see is low and there is a need to keep long-term statistics on the router. For example, if the only key field in the flow record is the 8-bit IP type of service (ToS) field, only 256 flows can be monitored. To monitor the long-term usage of the IP ToS field in the network traffic, a permanent cache can be used. Permanent caches are useful for billing applications and for an edge-to-edge traffic matrix for a fixed set of flows that are being tracked. Update messages will be sent periodically to any flow exporters configured according to the "timeout update" setting.

Introduction of New Show Commands

This new flexibility comes with new, powerful, show commands. For instance, the "top talkers" feature has been revamped in a new way and helps you analyze the large amount of data that Flexible NetFlow captures from the traffic in your network by providing the ability to filter, aggregate, and sort the data in the Flexible NetFlow cache as you display it.

The following example combines filtering, aggregation, collecting additional field data, sorting the flow monitor cache data, and limiting the display output to a specific number of high-volume flows (top talkers). It lists the top four (in terms of bytes transferred), aggregated by IPv4 destination address, filtered to match only protocol 1 or 6 (respectively Internet Control Message Protocol [ICMP] and TCP):

```
Router# show flow monitor FLOW-MONITOR-1 cache filter ipv4 protocol regexp (1|6)
aggregate ipv4 destination address collect ipv4 protocol sort counter bytes top 4
```

```
Processed 26 flows
Matched 26 flows
Aggregated to 13 flows
Showing the top 4 flows
```

IPV4 DST ADDR	flows	bytes	pkts
=====	=====	=====	=====
172.16.10.2	12	1358370	6708
172.16.10.19	2	44640	1116
172.16.10.20	2	44640	1116
172.16.10.4	1	22360	559

For more details, see "Using Cisco IOS Flexible NetFlow Top N Talkers to Analyze Network Traffic."

No SNMP Support

Traditional NetFlow did provide some Simple Network Management Protocol (SNMP) support, most notably to configure Traditional NetFlow and do some limited data polling such as the top-talkers table.

At this stage, Flexible NetFlow doesn't support SNMP configuration or data polling. Although the only current way to configure Flexible NetFlow is through the CLI, we are actively participating in development of a standard configuration model. All flow data may be exported to NetFlow collectors.

The direct effect is that tools used to automatically configure NetFlow using SNMP will not work with Flexible NetFlow.

Front-End Management

New data export, new collectors, new flows exported, new aggregation mechanism: all those changes open new possibilities, and that means updating your front end to support Flexible NetFlow.

Here are some applications that support NetFlow Version 9 to some extent. Due to the very nature of the IT sector, this list might change at any time and is certainly not exhaustive, but it gives you some pointers:

- Cisco NetFlow Collector (<http://www.cisco.com/en/US/products/sw/netmgts/ps1964/>)
- flowd (<http://www.mindrot.org/projects/flowd/>)
- Scrutinizer (http://www.plixer.com/products/scrutinizer_activate-netflow.php)
- SolarWinds Orion NTA (<http://www.solarwinds.com/products/orion/nta/>)
- CA eHealth Network Performance Manager (<http://www.ca.com/us/network-performance.aspx>)
- Java NetFlow Collect-Analyzer (<http://www.networkuptime.com/tools/netflow/jnca.html>)
- AdventNet NetFlow Analyzer (<http://www.manageengine.com/products/netflow/>)

Flexible NetFlow Migration in Practice

The Super Easy Way

A user can configure both Traditional NetFlow and Flexible NetFlow on an interface at the same time, and neither feature will have knowledge of the other. It is, however, recommended that this configuration be avoided as it might consume substantial resources.

As this is a trivial case that does not really use Flexible NetFlow, we won't talk about that there, but this might be an option you should be aware of.

Quick Jump from Traditional to Flexible NetFlow

Flexible NetFlow includes several predefined records that you can use right away to start monitoring traffic in your network. The predefined records are available to help you quickly deploy Flexible NetFlow.

If you have been using original NetFlow or original NetFlow with aggregation caches, you can easily continue to capture the same traffic data for analysis when you migrate to Flexible NetFlow by using the predefined records available with Flexible NetFlow.

Flexible NetFlow predefined records are based on the original NetFlow ingress and egress caches and the aggregation caches. Many users will find that the preexisting Flexible NetFlow records are suitable for the majority of their traffic analysis requirements. Thanks to predefined records, the migration from Traditional NetFlow to Flexible NetFlow is transparent to the collector and does not require the collector to be touched.

The difference between the original NetFlow aggregation caches and the corresponding predefined Flexible NetFlow records is that the predefined records do not perform aggregation. This is an advantage in that when someone only needs four NetFlow fields to track application usage, one can simply track those four key fields in Flexible NetFlow and the aggregation is natural.

This is contrasted to seven fields in Traditional NetFlow. In Traditional NetFlow, the user must track the seven key fields, and each field tracked leads to a greater number of flows that must then be aggregated.

Note: The difference is when the Cisco IOS Traditional NetFlow Aggregation feature is in use. In this case, Cisco Traditional NetFlow will summarize NetFlow export data on a Cisco IOS Software router before the data is exported to a NetFlow data collection system. The corresponding Flexible NetFlow predefined records do not perform aggregation, because it is implicit in the definition of the flows to track.

Flexible NetFlow predefined records are associated with a Flexible NetFlow flow monitor the same way as a user-defined (custom) record.

Let's convert this Traditional NetFlow sample to Flexible NetFlow:

```
interface FastEthernet 0/1
  ip flow [ingress|egress]
  exit
ip flow-export destination 192.168.9.101 9996
ip flow-export source FastEthernet 0/1
ip flow-export version 5
ip flow-cache timeout active 1
ip flow-cache timeout inactive 15
```

With Flexible NetFlow:

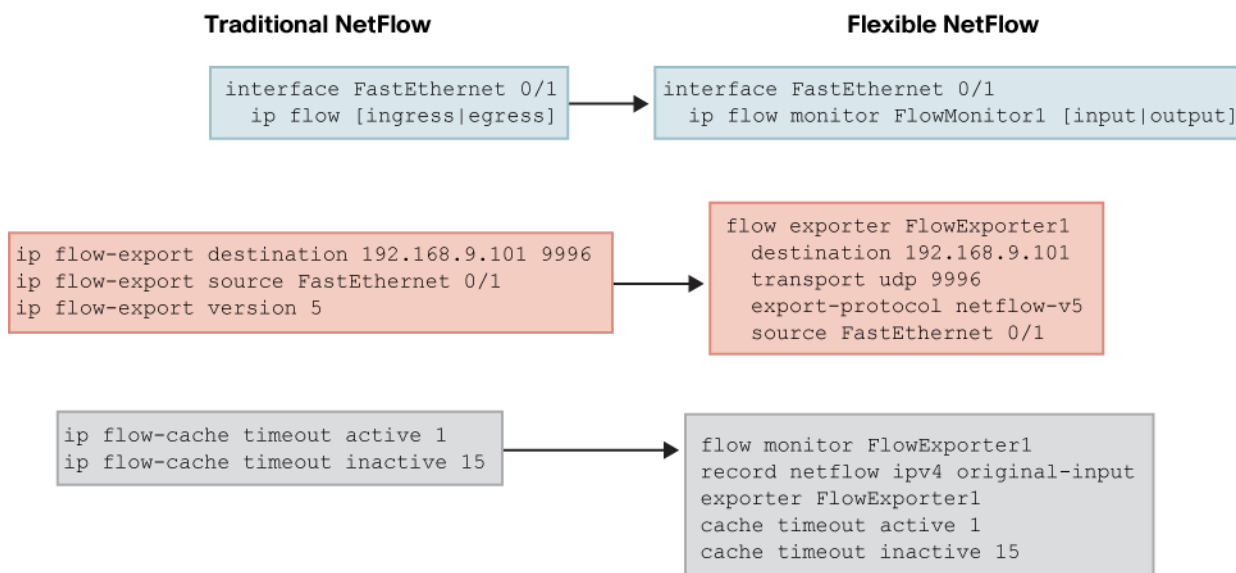
```
flow exporter FlowExporter1
destination 192.168.9.101
transport udp 9996
export-protocol netflow-v5
source FastEthernet 0/1

flow monitor FlowMonitor1
record netflow ipv4 original-input
exporter FlowExporter1
cache timeout active 1
cache timeout inactive 15

interface FastEthernet 0/1
ip flow monitor FlowMonitor1 [input|output]
```

A different way to present this modification is by illustrating the different components that used to be bundled together in Traditional NetFlow and are now separate entities in Flexible NetFlow: flow monitor, flow exporter, and interface. Note: If, unlike here, you don't use a predefined record, you'll also have a "flow record" configured. (See Figure 2.)

Figure 2. Configuration Sample in Traditional NetFlow versus Flexible NetFlow



Note: In some versions of Cisco IOS Software the "ip flow ingress" is the equivalent command for "ip route-cache flow."

Using the Predefined Records

Now that you've seen the basics, you might wonder how to translate your very own configuration if it does not exactly match the previous example. There are many predefined records to make your transition to Flexible NetFlow easy and painless.

So far, we have made these predefined records available:

- Flexible NetFlow "NetFlow Original" and "NetFlow IPv4 Original Input"
- Flexible NetFlow "NetFlow IPv4 Original Output"
- Flexible NetFlow "NetFlow IPv6 Original Input"
- Flexible NetFlow "NetFlow IPv6 Original Output"
- Flexible NetFlow "Autonomous System"
- Flexible NetFlow "Autonomous System ToS"
- Flexible NetFlow "BGP Next-Hop"
- Flexible NetFlow "BGP Next-Hop ToS"
- Flexible NetFlow "Destination Prefix"
- Flexible NetFlow "Destination Prefix ToS"
- Flexible NetFlow "Prefix"
- Flexible NetFlow "Prefix Port"
- Flexible NetFlow "Prefix ToS"
- Flexible NetFlow "Protocol Port"

- Flexible NetFlow "Protocol Port ToS"
- Flexible NetFlow "Source Prefix"
- Flexible NetFlow "Source Prefix ToS"

For more information and details for each of these predefined records, including descriptions of the key and non-key fields, please see ["Configuring Cisco IOS Flexible NetFlow with Predefined Records."](#)

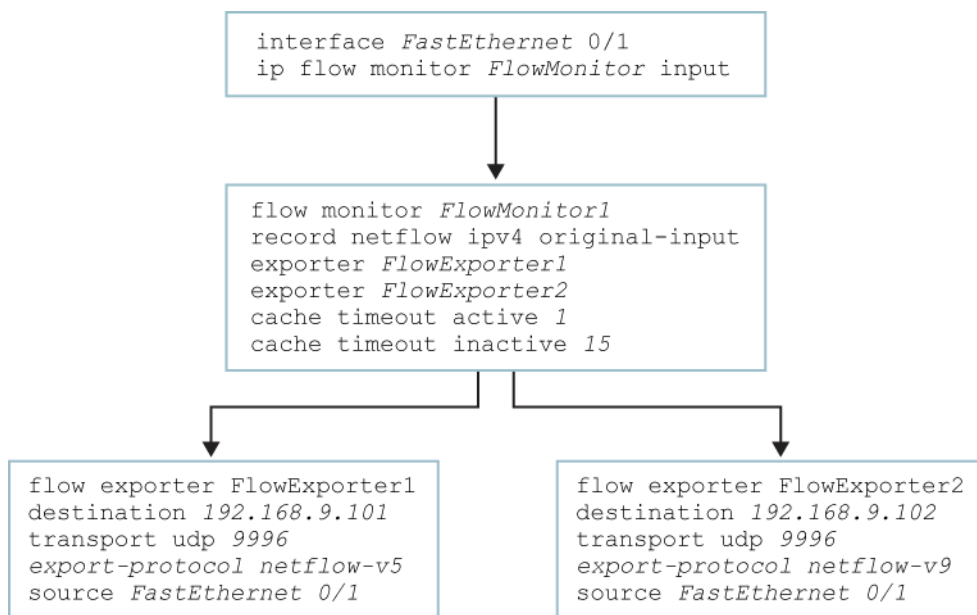
Next Steps: Using Flexible NetFlow for Real

Once you have Flexible NetFlow running in "backward compatibility mode" with Traditional NetFlow, you are safe. Your traffic is monitored, it is exported to your existing NetFlow collector, and everything goes well. You can already feel the new flexibility just by looking at the new show commands.

Now is a good time to go one step further: you can export exactly the same flows to both v5 and v9 collectors at the same time and start playing with your new collector without disturbing your existing infrastructure.

Let's start from an earlier example and have two collectors (v5 and v9) to the same flow monitor that you just migrated to Flexible NetFlow (Figure 3).

Figure 3. Hierarchy between Interface, Flow Monitor, and Flow Exporter



Once you have a working NetFlow Version 9 collector, you can move one step further and set two flow monitors on the same interface: one that emulates Traditional NetFlow and exports in v5, and one that is really unleashing Flexible NetFlow and exports with v9 and uses your own flow record.

Let's see what our configuration would look like:

```
flow exporter FlowExporterTrad
destination 192.168.9.101
transport udp 9996
export-protocol netflow-v5
source FastEthernet 0/1

flow exporter FlowExporterFlex
destination 192.168.9.102
transport udp 9996
export-protocol netflow-v9
source FastEthernet 0/1

flow monitor FlowMonitorTrad
record netflow ipv4 original-input
exporter FlowExporterTrad

flow record FlowRecordFlex
match ipv4 section payload size 900
match transport udp destination-port
match ipv4 destination address
match ipv4 source address
collect counter packets

flow monitor FlowMonitorFlex
record FlowRecordFlex
cache type immediate
cache entries 1000
exporter FlowExporterFlex

interface FastEthernet 0/1
ip flow monitor FlowMonitorTrad input
ip flow monitor FlowMonitorFlex input
```

Conclusion

We have made everything possible to help you make the transition from your current Traditional NetFlow environment. Tools such as predefined records emulating the Traditional NetFlow you know today will speed up Flexible NetFlow adoption.

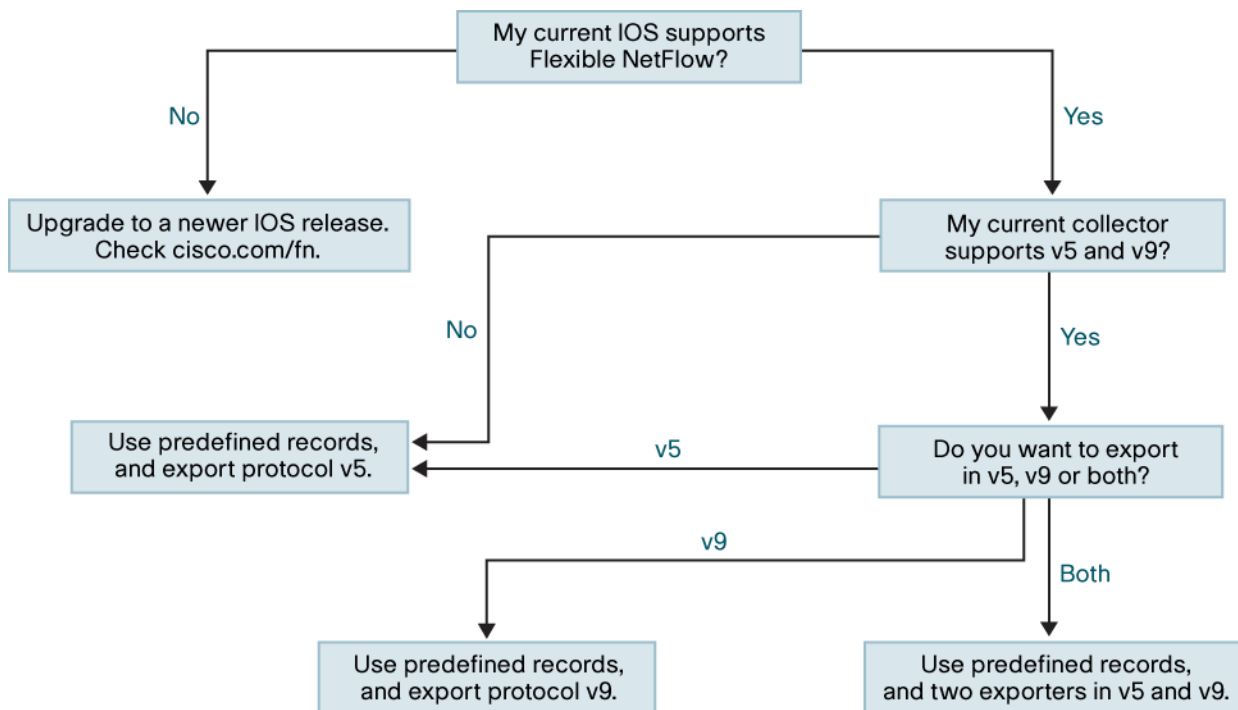
While you work in "backward compatibility mode," you know that you're safe, and you can start exploring Flexible NetFlow, including the show commands, having another collector running v9, defining a new flow monitor with a different aggregator, or using multiple collectors with different cache types. You can do all that while preserving the security and comfort of emulating Traditional NetFlow.

And when you feel comfortable, make the switch.

Welcome to Flexible NetFlow.

Figure 4 shows a recommendation based on your current system and wishes.

Figure 4. Decision Chart for a Painless Migration to Flexible NetFlow



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCI, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)