



NetFlow

Peter J. Welcher

Introduction

I'm excited, I'm just about done preparing for our first teach of the Introduction to MPLS course, MPLS Essentials (MLSTE). Our first public class is 5/30/2001, and we've scheduled additional classes in July and August.

Congratulations to Mentor Technologies' Scott Morris, now double-CCIE in Routing & Switching, and ISP Dial, with Security Specialization. (But, Scott, if you pass any more certs your signature will cause buffer overflows!) Scott teaches CIT, ICCR (cable), and soon CIPT (IP Telephony/Call Manager).

For useful links, prior articles, etc., see <http://www.netcraftsmen.net/welcher/index.html>.

For some consulting/training work I'm doing, I've been reviewing, reading about, and working with NetFlow, NetFlow FlowCollector, and NetFlow Data Analyzer. Although I've mentioned NetFlow in passing in other articles, I hope that describing all of these in one article may be useful to you.

What is NetFlow and Why Do I Care?

NetFlow has meant several things to Cisco customers at various times. I prefer to understand this historically: NetFlow evolved as a caching technique. To speed up network flows (source IP, source port, destination IP, destination port) and Layer 3 switching in the presence of access lists, the Cisco router and switch caches were re-organized based on the flow information. As this code became more efficient, a side benefit was the collection of useful flow statistics, without too severe a performance penalty. Even with CEF (Cisco Express Forwarding) for rapid Layer 3 switching, NetFlow caching can apparently still enhance performance of longer access lists (more than 10 to 25 entries or so), Policy Routing, and perhaps other features ("NetFlow feature acceleration"). But there is also real benefit to the reporting data it provides.

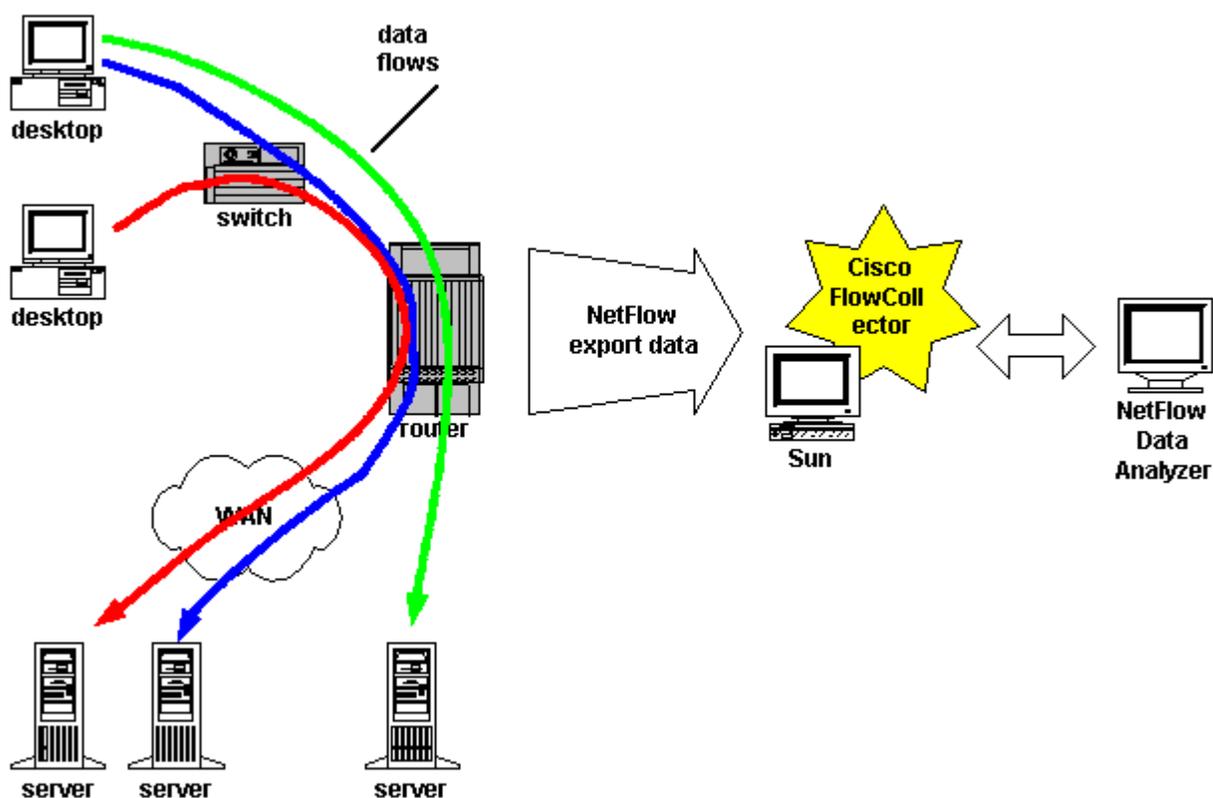
So there are two reasons you might be using NetFlow: to speed certain access list uses up, or to collect data. If you're only interested in speeding up, you can skip all the words here about reporting and jump right to the part about how to configure NetFlow ([see below](#)).

Typical uses for this data: tracking what kind of traffic is entering or exiting an ISP or corporate network, tracking traffic flows between BGP Autonomous Systems (AS's), or enterprise network regions, etc. Lately several third parties are providing billing software, so that IP Service Providers can bill customers for the data sent or bandwidth used. For example, Digital Island's billing is based on

NetFlow data, charging per Megabyte of data delivered in-country to in-country ISP (see also <http://www.digitalisland.com> , for example <http://www.digitalisland.net/news/press/lcm2.shtml>).

The way that NetFlow reports statistics is by flow export. As cache entries expire (or are actively expired by an algorithm), the packet/byte count data for the unidirectional flow is exported to a collector station. This would typically be a Sun workstation running the Cisco FlowCollector software, although a company called Apogee provides NT-based NetFlow collection and reporting software.

The Cisco FlowCollector software aggregates the data, rolling up raw data according to the aggregation scheme or schemes you have selected ([see below](#)). The Cisco NetFlow Data Analyzer software then provides some viewing, searching, and graphing of the aggregated data. If you have enough FlowCollector stations, you might instead roll the data up to a data warehouse, and run your billing or reporting off of that. (If you're doing that, you probably will have some consultants in doing the system integration work. At larger scales, this is very noticeably not shrinkwrap software.)



You need to understand this about aggregation: it still is flow-related information. It will tell you totals per flow, or per source address or destination address or per interface/source/destination combination. If you want totals per interface (utilization), you have to add the numbers up for yourself. This may not be that hard, but it is **not** part of what FlowCollector and Data Analyzer do for you. On the other hand, they're great for examining all the individual flow contributions to a busy link!

For the GSR 12000 router, volume of NetFlow export is apparently a potential CPU-related issue. One Cisco answer is the on-board router-based aggregation. Another is statistical sampling, where NetFlow only collects statistics on a fraction of the packets traversing the router.

Another simple approach that became available with Cisco IOS 12.1(5)T is Traffic Matrix Statistics,

TMS. I've been thinking of TMS as simplified NetFlow for backbone BGP speaking routers. It measures "internal" and "external" traffic. See

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/tms.htm> .

What Sort of Reports Can I Get From NetFlow?

If you've made it this far, you're probably wondering, well, what kinds of report can I get out of this system?

Here's the short form version of the various aggregation schemes. The table is followed by the link to the online Cisco documentation, which will give you full details. All the aggregation schemes report on packet counts, byte counts and flow counts. The difference lies in the key fields, that is, how much detail as to source, destination, port, ToS byte, etc. is provided.

Name of Aggregation Scheme	Key Fields
RawFlows	All
SourceNode	Source address
DestNode	Destination address
HostMatrix	Source and destination addresses
SourcePort	Source port
DestPort	Destination port
Protocol	IP sub-protocol
DetailSourceNode	Source address, source and destination port, protocol
DetailDestNode	Destination address, source and destination ports, protocol
DetailHostMatrix	Source and destination addresses and ports, protocol
DetailInterface	Source and destination addresses, input and output interfaces, next hop address
CallRecord	Source and destination addresses and ports, protocol, ToS
ASMatrix	Source AS and destination AS
NetMatrix	Input and output interfaces, masked source and destination address, source and destination masks
DetailASMatrix	Source and destination address and port, protocol, input and output interfaces, source and destination AS
ASHostMatrix	Source and destination address, source and destination AS
HostMatrixInterface	Source and destination address, protocol, input and output interface
DetailCallRecord	Source and destination address, source and destination port, protocol, ToS, input and output interface

Quibbles:

- 1 The ToS byte allows you to report data based on IP Precedence or DSCP, which is a step towards managing QoS.

- 1 The AS info can either be the AS of the source and destination, or it can be the AS of the peer reached through the input or output interface.

For full information about the various aggregation schemes, see the documentation chapter, Customizing FlowCollector, specifically the part beginning at

http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/nfc/nfc_3_0/nfc_ug/nfctune.htm#xtocid93571

A relatively new feature in Cisco IOS 12.0(3) T allowed you to configure aggregation within the router. When you do this, the router exports version 8 record instead of the usual version 5 records. (Think of "version" here as more like "type of record" exported. You'll see version 5 normally, and 7 if you have Catalyst 5000 switches with NFFC.) FlowCollector 3.0 understands the version 8 records, and NetFlow Data Analyzer 3.0 can report on them. For more about this new Cisco IOS feature, see

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t3/netflow.htm> .

Another newer feature allows you to specify the granularity of addresses, by specifying a minimum mask that all source and destination aggregate addresses are logically ANDed with. (For example, 24 bits means you see /24 or shorter masks, or no more than 24 bits of subnetting.) Different minimum masks can be set for source and destination address. See

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t2/dtnfmask.htm> .

The following table shows the router-based aggregation schemes that are now available:

Name of Aggregation Scheme	Key Fields
RouterAS	Inbound and outbound interfaces, source and destination AS
RouterProtoPort	Source and destination port, protocol
RouterSrcPrefix	Input interface, source AS, masked source address, source mask
RouterDstPrefix	Output interface, destination AS, masked destination address, destination mask
RouterPrefix	Input and output interface, source and destination AS, masked source and destination addresses, source and destination masks

One big advantage to aggregating on the router is that it can greatly reduce the volume of data exported. This saves router CPU and bandwidth, and also saves disk space on the collector station. However, if you are exporting version 5 "raw" NetFlow to the Collector station, adding router aggregation just makes the router work harder, and it would probably be better to just aggregate on the collector. In other words, decide in advance where to aggregate. If you need raw data anyway, or if you don't like the router aggregation schemes, aggregate on the collector station. But if you can, aggregate on the router.

What is NetFlow FlowCollector?

NetFlow data from older routers is version 1 format. Version 5 is more recent. Version 7 comes from NFFC in Catalyst 5000. And Version 8 comes from router performing router aggregation.

NetFlow FlowCollector understands how to receive all these, aggregate them into various buckets depending on how you've configured it, filter out unwanted data, and store to either ASCII or compressed binary file. (Storage is no small thing, NetFlow can fill up your disk in a hurry!) The storage scheme is heirarchical, to help you find your data later. And FlowCollector also helps manage file

system space.

FlowCollector has four (4) subsystems: Collector, Gateway, Daemon, User Interface. The Collector collects the data, the Daemon makes sure the Collector is running, the User Interface is an ASCII menuing scheme for configuring the Collector, and the Gateway talks to other programs, for example the NetFlow Data Analyzer. The preferred way to configure and control one or multiple FlowCollectors is through the NetFlow Data Analyzer GUI.

Cisco FlowCollector overview:

http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/nfc/nfc_3_0/nfc_ug/nfcover.htm

In summary, Cisco FlowCollector is (or should be) the silent partner in all this. It is the set of background workhorse programs that actually collect and filter your data, and make it available to the Data Analyzer and other User Interface tools. The threads do the job of aggregation. You do need to be selective about what aggregation schemes you use, since there is a limit of 10 active threads (aggregation schemes) per Collector station.

What is NetFlow Data Analyzer?

Cisco NetFlow Data Analyzer (NDA) overview:

http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/nda/nda_iug/nda_over.htm

NDA provides a GUI for retrieving, searching, viewing, and graphing NetFlow and TMS data. It also provides a simple GUI to help you set up (configure) NetFlow on your routers. (This GUI won't configure routers requiring a username login.) The simple GUI also allows you to control the Collector processes on multiple collection stations, including activating and de-activating threads, specifying ports to listen to, and specifying filters and aggregation schemes. You can pull in flow data by date and time, and show the top 100 or 500 or more flows (fixed sizes only, up to 10000 top flows -- and you really don't **want** to be looking at more!) Then you can sort by any of the columns. You can pick any column and row and histogram that flow: plot the time values of the specified variable. So you can see when the flow has bursts. (Admittedly, the X-axis labelling is atrocious, but you can position the cursor near data points for readable values.) What you can't do is multi-select rows and histogram them (plot the time series utilization). Nor can you use a pattern to select all the data meeting match criteria, and then aggregate it again on the spot. (Which would be very convenient). You can export to a CSV file to load into Excel. That's how you're going to be printing, by the way. To print graphs, use screen capture. Shall we say, this is currently interesting, useful, but perhaps unpolished software?

Tip: using the DetailedASMatrix aggregation scheme allows you to use the AS Drill Down feature in Analyzer, which is rather useful. If you're an enterprise network running no BGP, don't forget that you can just use AS 0 for the reports (0 is what it'll be recorded as, if you don't configure one of the AS option for your flow export, see below).

Configuring NetFlow and NetFlow Export

- 1 NetFlow Switching Overview:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnf
- 1 For a high level guide to configuring NetFlow, complete with examples, see the online Cisco NetFlow Configuration Guide:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_c/swprt2/xcfnf

- 1 For the details of syntax, defaults, etc. for each of the NetFlow commands discussed below or in the Cisco Configuration Guide, see the Cisco IOS Switching Services Reference Guide: http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fswtch_r/index.htm

Sample configuration enabling NetFlow on an interface:

```
interface fastethernet 0/0
ip route-cache flow
```

On VIP, add `ip route-cache distributed` before configuring NetFlow.

To tell the router where to send the NetFlow data, enter the following global configuration command:

```
ip flow-export ip-address udp-port [origin-as | peer-as]
ip flow-export version 5
ip flow-export source loopback 0
```

The `ip-address` here is the FlowCollector station's address, and the `udp-port` is the UDP port the FlowCollector thread is listening to (defaults are 9995 and 9996). Use the **origin-as** keyword for the AS of the source or destination address. Use **peer-as** for the AS of the ISP peer used to reach the source and destination. Omit either keyword if you have a non-BGP-speaking router (to save a few CPU cycles).

Specify **version 5** or some of the recent reports (aggregation schemes) won't work. You might want to specify the source address for the flow export, or else when routing shifts, FlowCollector will see the new source address (outgoing interface the flow records are sent from) as a different address, hence a different router. (You can create a RouterGroup in Analyzer to work around this, but that's a nuisance.)

You can tweak the size of the NetFlow cache, or break your router if you get carried away, by configuring:

```
ip flow-cache entries number
```

The number of entries can be from 1024 to 524288. The default is 65536.

EXEC mode commands for NetFlow:

```
show ip route flow
show ip flow export
clear ip flow stats
```

To configure router aggregation, globally configure

```
ip flow-aggregation cache aggregation-scheme
```

where *aggregation-scheme* is one of: **as**, **destination-prefix**, **prefix**, **protocol-port**, or **source-prefix**. This puts you into flow-cache configuration mode. You can then enter commands such as:

```
Router(config-flow-cache)#cache entries 2046
Router(config-flow-cache)#cache timeout inactive 240
Router(config-flow-cache)#cache timeout active 45
```

```
Router(config-flow-cache)#export destination 100.1.2.3 9990
Router(config-flow-cache)#enabled
```

The first specifies the number of cache entries for the aggregation scheme. The second entry is the timeout for an inactive flow, in **seconds**. The third specifies the timeout for an active flow, in **minutes**. The fourth is where to send the version 8 router aggregated flow records (ip address of collector station, and UDP port). Finally, we enable the aggregation.

To configure the new minimum mask feature (with appropriate router aggregation scheme), add as appropriate:

```
Router(config-flow-cache)#mask source minimum value
Router(config-flow-cache)#mask destination minimum value
```

EXEC mode command relating to router flow aggregation:

```
show ip cache flow aggregation
```

See the documentation for more details or any commands I've missed.

In Conclusion

Cisco main NetFlow link: <http://www.cisco.com/warp/public/732/netflow/>

Cisco NetFlow Partners link: <http://cisco.com/warp/customer/732/partners/nfpartner.html>

Cisco link for DEMO trial copies of NetFlow FlowCollector and Data Analyzer software: <http://www.cisco.com/kobayashi/sw-center/netmgmt/netflow/nf-planner.shtml>

NetFlow-related freeware or other links:

- 1 CFLOWD: <http://www.caida.org/tools/measurement/cflowd/>
- 1 NetFlowMet: <http://www.auckland.ac.nz/net/NeTraMet/>
- 1 Smurfind: ftp venera.isi.edu, subdirectory mon. The exact file names are README.smurfind, smurfind.c, flowdata.h. I tried to connect to test this a couple of times and couldn't (5/8/01). See also <http://www.cctec.com/maillists/nanog/historical/9901/msg00425.html>.
- 1 Extremely Happy NetFlow: <http://ehnt.sourceforge.net/>

Just scanning the links on NetFlow found with a quick Web search, it looks like NetFlow is becoming an ad hoc Internet reporting standard. If you know of other vendors that support NetFlow, or if you have some links you'd like to share through me, please email them.

Your comments, preferences and ideas and suggestions for topics are always more than welcome! I enjoy hearing from you!

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner dedicated to quality consulting and knowledge transfer. NetCraftsmen has nine CCIE's, with expertise including

large network high-availability routing/switching and design, VoIP, QoS, MPLS, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher> . New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to pjw@netcraftsmen.net .

5/8/2001

Copyright (C) 2001, Peter J. Welcher