

5 Tips for Improving NetFlow Monitoring Performance

Cisco's IOS NetFlow service collects statistical data about IP-based traffic flows for many purposes, including:

- **Network monitoring.** Traffic flow analysis reveals traffic patterns across the network and at the element or interface level. Monitoring these patterns, and the sources of IP packet flows, can help detect and troubleshoot performance and availability issues.
- **Application monitoring.** Monitoring application bandwidth shows the network load associated with new and existing applications and identifies opportunities for reallocating resources to improve application performance.
- **User monitoring.** Understanding how users consume network resources provides important usage and planning information, as well as information regarding security and policy violations.
- **Security monitoring.** Detecting intrusions such as denial-of-service attacks, viruses, and worms in near real-time is critical for maintaining the integrity of the network.
- **Network/capacity planning.** Capturing data over time provides the information needed to plan network improvements and increase capacity such as additional routers or T1 lines.

NetFlow can be a powerful tool... when you need a deeper look into traffic traveling through your network. Many network engineers are reluctant, unfortunately, to use NetFlow due to fears—often with good reason—NetFlow data collection might put performance burdens on the network. Although you may have heard that NetFlow is a performance hog, with proper tuning, you can usually get the data you need in an acceptable way.

Most enterprise-level IT groups use NetFlow (or similar protocols) to collect flow data. Administrators in smaller organizations, however, worry about NetFlow's impact on network performance. In truth, NetFlow can be a resource hog if used indiscriminately. With proper planning, however, you can use NetFlow statistics to get valuable planning, monitoring, and troubleshooting data at an acceptable performance cost.

Netflow uses router CPU resources proportionally according to the number of IP flows that pass through the router's inbound port (called "ingress" in Cisco-speak). The more IP flows present, the more resources NetFlow requires. Here are five tips for reducing the impact on router and network resources when collecting IP flow data with NetFlow.

1 Activate NetFlow on Edge Routers

Avoid collecting duplicate flow data by activating NetFlow only on edge routers where traffic originates or terminates. Activating NetFlow on backbone routers or intermediate routers is an unnecessary waste of valuable router and network resources given the flows were already sampled and exported by an edge router. Also, reduce NetFlow bandwidth requirements by locating the NetFlow collector(s) as close to the edge router as possible, ideally on the same switch as the router.

2 Use Input Filters

NetFlow offers hundreds of data points for dozens of uses. The best way to limit NetFlow's performance impact is to know what you want to achieve, then use input filters to limit incoming traffic to the things you really care about. For example, if you're only interested in accounting applications, then you likely only require basic router flow information. Network monitoring, on the other hand, may require more data, such as protocol and service types. Filters can classify packets according IP source and destination addresses, Layer 4 protocol and port numbers, incoming interface, MAC address, and so on. Using input filters restricts the volume of flows being sampled, which reduces NetFlow's affect on router CPU utilization.

3 Apply Sampling Rates

By default, NetFlow creates flow records for all inbound packets, which can significantly increase router CPU utilization. You can reduce this affect, however, by sampling every 100 or 1000 packets rather than all packets. Table 1 shows the difference in CPU utilization for a Cisco 12000 router employing 1:100 sampling versus full sampling and compared to CPU utilization without NetFlow.

Number of Flows	10K	45K	65K
No NetFlow sampling	7%	8%	11%
1:100 NetFlow sampling	9%	12%	14%
Full NetFlow sampling	12%	23%	31%

Table 1: Differences in CPU Utilization Rates for a Cisco 12000 Router. Differences will vary depending on router model.

As you can see, 1:100 sampling results in only a small increase in CPU utilization. Limited sampling can be very practical for capacity planning or network planning tasks, for example, since data for every flow is not needed. You can set the sampling rate as needed to anything from 1:1 to 1:65535 depending upon the granularity of data required for your objective.

When configuring a sample map, you can enable the following sample types:

- **Deterministic sampling** selects every nth packet.
- **Random sampling** selects one out of every n packets.
- **Time-based sampling** selects a packet every n milliseconds.

In most cases, random sampling will produce the most representative flow statistics.

In addition, combining input filters with sample maps can help balance the need for data against performance cost. For example, you could create filters for high-, medium- and low-priority traffic flows, then configure sample maps with 1:1, 1:100, and 1:1000 sampling, respectively. This provides full sampling for a small but critical subset of flows, with broader sampling for other flows, all at a reasonable cost.

4 Optimize Cache Aging

When a flow “expires,” or ends, its flow record is exported from the NetFlow cache to one or more NetFlow collectors. By default flows that exceed 30 minutes are considered expired, even if the flows are still active. Because exporting flow records consumes CPU utilization, you can reduce the load by increasing the interval at which long-lived flows are allowed “age” or remain in the cache. To optimize cache aging, try setting the Active Time value to 40 or 45 minutes (the maximum value is 60 minutes).

5 Enable Data Aggregation

You can reduce both the bandwidth requirements for exporting NetFlow data and the system requirements for NetFlow collectors by enabling router-based aggregation. NetFlow Export Version 9 (v9) provides 11 router-based aggregation schemes for structuring data prior to export, such as source prefix, destination prefix, and protocol-port aggregation. Note that there is little or no router impact when enabling v9 aggregation.

Balancing the need for relevant data at an acceptable performance cost is the primary goal of any NetFlow deployment. Careful planning before you enable NetFlow can help you achieve this goal in short order. For guidelines on how to plan and implement your NetFlow deployment, see [Cisco's NetFlow Services Solutions Guide](#).

Collecting Netflow Data with dopplerVUE

With dopplerVUE, you can analyze NetFlow data out of the box. There are no add-on modules or plug-ins, and no additional cost. Simply configure the router to send NetFlow data to the dopplerVUE IP address and you're ready to go.

dopplerVUE allows you to graphically view 14 NetFlow metrics and generate reports showing traffic flow statistics at the device and interface level. Since NetFlow support is built in, graphs and charts can be combined in a single view alongside other relevant information, such as SNMP-based data about bandwidth usage, WMI data about server performance, as well as IP SLA, ICMP or other types of management data.



Figure 1: dopplerVUE's NetFlow charts show the nature of traffic flow across your network.

About dopplerVUE

dopplerVUE is a powerful yet easy to use network management solution for managing up to 5000 network elements. To see how dopplerVUE can solve your network management needs, visit to www.dopplerVUE.com. For more information call 888-dvue-now (888-388-3669).

© 2008 SYS Technologies. All Rights Reserved. All trade names referenced are the service mark, trademark or registered trademark of the respective manufacturer.