



Purpose

This document describes what network flows are, and also compares and contrasts the multitude of different types of flows.

Target Audience

The document is targeted at network designers who need more information on available technologies for monitoring their environment.

Flow Definition

A network flow is a relatively new term, or at least an obscure definition. When Cisco originally released the Netflow protocol, they defined a source to be a series of packets that share all of these values:

- Source IP Address
- Destination IP Address
- IP Protocol
- Source Port
- Destination Port
- Ingress Interface
- IP ToS

A series of packets that shares all of these qualities is classified as a flow, and is typically saved on the routing device/network observer until a certain timeout is met or until the device's cache is filled, at which point, the device's flow data is forwarded to a flow collector, like Nagios Network Analyzer for evaluation and analysis. Information that is kept about these flows is quite exhaustive and contains much of the information that you would like to know about the actual communication going on on your network.

Implementation Equivalency

There are a number of flow protocols: Netflow, Jflow, Cflow, sFlow, etc. This can be very confusing for an admin or designer who is just getting started. The bottom line is, most of them are at least superficially equivalent to Netflow. The one protocol that is absolutely not functionally equivalent of Netflow, is sFlow.

sFlow takes a different approach to network traffic analysis, it was developed with an eye towards scalability. Netflow data samples every packet that passes through an interface. sFlow, on the other hand, takes samples of the packets coming through, and uses some statistical analysis to make guesses about the flow data that is flowing on the network. For instance, a saturated 10 gigabit interface is going to generate an insane amount of flow data when using Netflow and can bog down the CPU and generate even more network traffic to send its flow reports. sFlow, can take significantly less resources to monitor the same resources, at the expense of loss of granularity and possibly missing some connections altogether. It is a give and take decision.

Flow Versions

The actual information that is kept in flows varies from implementation to implementation, and even if you're using the Netflow standard, the information carried can vary from version to version! If you are deeply interested in Netflow versions, then you should probably read the following link:

http://netflow.caligare.com/netflow_format.htm

IPFIX is closely associated with Netflow versioning, to clear up any ambiguity, this is an excellent resource:

http://www.ist-lobster.org/events/tutorial-2005/tutorial_netflow.pdf

However, if you are more interested in a simple cursory explanation of the differences, perhaps the following will suffice (please note this is not anywhere close to exhaustive, but more shows some milestones for general use):

- v5 - The baseline version, will be lowest that will be included on a device, does not support Ipv6
- v9 - Added support for Ipv6, and flexibility for field types (user definable)

Finishing Up

This completes the documentation on understanding network flows in Nagios Network Analyzer.

If you have additional questions or other support related questions, please visit us at our Nagios Support Forums:

<https://support.nagios.com/forum>

The Nagios Support Knowledgebase is also a great support resource:

<https://support.nagios.com/kb>