

Successful Ways to Use NetFlow

This paper outlines a few different ways to use NetFlow technology, as distributed traffic analysis is just the beginning. Understanding application behavior is another NetFlow frontier.

What is NetFlow?

NetFlow is a technology developed by Cisco Systems that is also supported by many other vendors. NetFlow is far easier to setup and use than other technologies that require the deployment of distributed hardware probes. Basically, the router or switch is setup to summarize every conversation incoming on each network interface. A single conversation is generally made up of several, or even hundreds, of packets. Up to 30 summarized conversations are compiled into an individual NetFlow packet.

NetFlow packets are then sent off to a NetFlow collector/analyzer. Although the bulk of the conversation detail is stripped out (i.e. information directly related to the application), the following is still preserved:

- Source and destination IP addresses
- Source and destination protocols
- Source and destination interfaces of the router/switch
- Type of Service (aka ToS)
- Number of packets
- Other details

Loaded with the above details, the collector can compile and provide incredible details on traffic behavior. Traditionally, distributed analysis has been limited to SNMP trending tools, such as MRTG, and packet analyzers, such as Ethereal (aka Wireshark). NetFlow provides the best of both tool sets, without physically running around the network.

How is NetFlow helpful?

Once you have NetFlow being exported to a collector, it is time to put that data to good use. The following examples outline a few ways to leverage the wealth of information exposed by NetFlow.

1st EXAMPLE

Which application uses the most bandwidth?

One way of leveraging NetFlow is using it to understand how much traffic a user creates with the typical utilization of an application. If the application is either console-based or web-based, the level of traffic can be dramatically different.

For example, a web-based application would not only download the data from a search, but must also download all of the HTML needed to generate the web page. This can include several drop down boxes for various menu options. This occurs each time a page is rendered!

Console based applications are typically installed on the end user's computer. Unlike web browser applications, they can become outdated. However, since the application is installed locally, the information needed to render a page is found locally, and only the database information is sent across the network. Even if the data volume isn't significant, performance generally is, as busy web servers can be the cause of seemingly slow applications.

To understand how much load a web application is putting on the network, a test should be performed:

- How much traffic does the console application generate when performing a function A, B and C?
- How much traffic does the web interface generate for the same application when performing the same functions A, B and C?

Equipped with the values above, this information can be loaded into equations which include the bandwidth of a WAN link and the number for users expected to use the application. In this scenario, a web-based application may become less attractive.

Many companies are now using thin clients (e.g. Citrix). These applications involve terminal servers and operate similar to main frame technology, where the user typically employs the local computer as a terminal into the system. These applications can come at a cost, when running over expensive lease lines.

2nd EXAMPLE

Rerouting lower priority traffic

Another way to benefit from NetFlow technology is by using it to understand internet traffic, like what needs to be on the corporate internet connection and what doesn't. Often, certain web sites can't simply be blocked!

Unfortunately, network administrators will need to dodge political upheavals pertaining to network access; and because of this, tools, like NetFlow analyzers, allow savvy admins to determine what network traffic can be pushed onto a backup internet connection, making room for legitimate, directly business related traffic on the company's primary, expensive internet connection.

If you are noticing excessive traffic to a host on a domain, such as ~.deploy.akamaitechnologies.com, then you know that traffic related to akamaitechnologies.com is related to ads in free mail utilities. Many of us like to check our non work related email during the work day.

Most administrators can't stop this. However, they may want the traffic this creates to get routed out a backup internet connection, in order to save bandwidth on the company's primary internet connection. Many routers have the ability to route traffic destined for specific hosts or subnets out an alternate interface.

3rd EXAMPLE

Identify problems without a packet analyzer

NetFlow can be leveraged similarly to a packet analyzer in order to ascertain "who is using all the bandwidth." Think of a NetFlow as a traffic trend with many more details than what an SNMP graph would provide.

Good NetFlow tools allow administrators to:

- Quickly navigate to specific interfaces.
- Zoom in on specific time frames.
- Click for top hosts, protocols and conversations specific to a time frame.
- Perform custom queries on existing captured data using multiple criteria, such as IP address and protocol.

Although, NetFlow may not provide the details offered by packet analyzers, most troubleshooting that traditionally required a packet analyzer can be resolved using NetFlow data. When troubleshooting an SNMP issue, for example, a packet analyzer displays the OID and community string used. NetFlow omits these details, providing only the source and destination IP address, protocol (i.e. SNMP), the amount of packets used and a few other details.

There are several NetFlow analysis tools available on the market. Johnny Espinosa from PSI used Scrutinizer NetFlow & sFlow Analyzer from plixer International, Inc. to troubleshoot an issue with failing nightly processes at one of his company's locations. Within a few days of implementing Scrutinizer's free evaluation software, Johnny was able to ascertain that the problem resided on PSI's FTP server. A process on the server was running nightly and saturating the WAN link, which in turn was causing other processes to fail from a lack of available bandwidth across the link. "After we were able to isolate the issue, we were able to setup QoS policies in place to accommodate for all processes. I believe if we had not had the Scrutinizer demo running, we would still be tracking this issue down. This is a great product and it's functionality has helped us mitigate issues before they become problems!"

Summary

NetFlow does not eliminate the need for packet analyzers; rather, it reduces the need for them. Further more, problems are more likely to be pursued, when using NetFlow, because the data is more readily available.