Whitepaper

NetFlow vs. sFlow:

A Technical Review

plixer International

Table of Contents

Abstract

In an effort to gain more insight into large scale networks where packet probes are not feasible, NetFlow and sFlow capable routers and switches are being utilized. NetFlow & sFlow are technologies supported by most major router and switch vendors whereby packet analyzer like details are pushed to a collector. This paper provides technical insight into the differences between the two.

Introduction

NetFlow Vs. sFlow is not so much a question of which is better, it is more of an architecture question of: Where should each be deployed? NetFlow (i.e. IP FIX) is a standard developed by Cisco and is generally software based. However, there are hardware implementations (e.g. Enterasys). Inmon is the developer of sFlow, which is hardware based.

NetFlow

When NetFlow version 5 is enabled on an interface, it caches conversations between hosts and exports the conversations in a configurable interval, which is typically every 60 seconds for TCP and immediately for UDP. The packets between host A and host B are summarized into a single record in a NetFlow datagram. A single NetFlow packet can contain up to 30 records where each represents potentially thousands of packets. Because of its aggregation method, it normally results in a less than a 1% increase in network traffic. Vendors supporting NetFlow can be found here:
http://www.plixer.com/products/scrutinizer_activate-netflow.php

sFlow

SFlow is a packet sampling technology. Some implementations can only sample every 100th packet per interface while others, such as Foundry, can sample every other packet. Although sFlow can provide more details than NetFlow, such as errors per interface, sFlow generally is not as accurate when measuring total traffic between two hosts. This is only true in pure IP environments. Vendors supporting sFlow can be found here:
http://www.sflow.org/products/network.php

Developments in NetFlow v9 allow it to sample similar to sFlow.

Lab Configuration

Hardware

In the lab, an Extreme Summit sFlow switch running v7.6 firmware was inserted between the Enterasys switch running Rev 05.42.04 and the firewall (SonicWall). The Enterasys switch supported NetFlow v9 and the Extreme switch supported sFlow v5. The sampling rate on the Extreme was configured to sample every packet. The lab technician wasn't confident that the Extreme Summit switch could sample every packet, but the switch didn't complain after entering the command.

Collection and Analysis

For flow collection, Scrutinizer NetFlow & sFlow Analyzer v6 was used, which is pictured below. PLXRSW3 (sFlow) is the Summit switch and PLXRSW1 (NetFlow) is the Enterasys Switch.

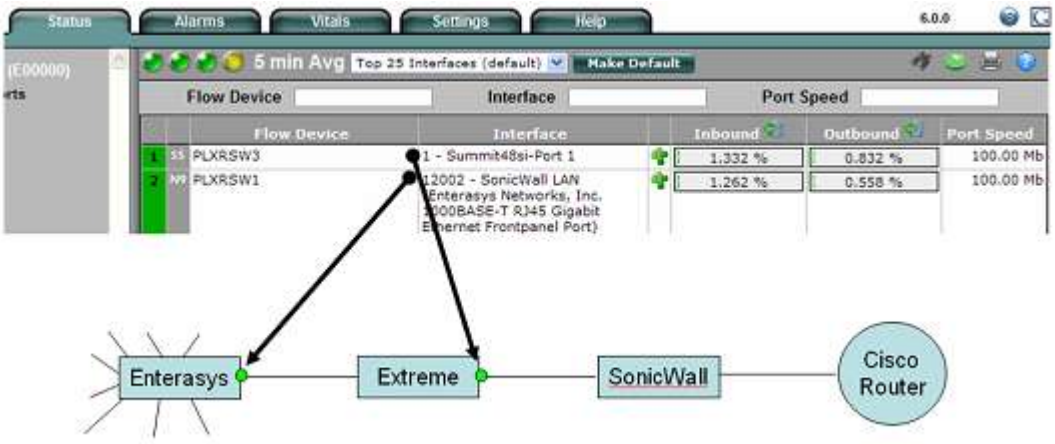

Utilization Measurements

The above configuration displayed traffic rates of the same live traffic using NetFlow and sFlow collection. Notice above that the Inbound and Outbound - five minute traffic averages don't match for exactly the same traffic volumes. The Extreme Summit = 1.332 % and the Enterasys = 1.262 % for Inbound utilization. The lab technician believes this was likely caused by many things including the fact that sFlow samples tend to be exported closer to real time. NetFlow, on the other hand, has to deal with active and inactive timeout configurations (http://www.plixer.com/products/scrutinizer_activate-netflow.php). Because of this, an sFlow switch would likely reflect a sudden spike in utilization quicker than a NetFlow switch.

At times both switches would be as much as 1% different from one another, but for the most part they were pretty much the same. Below is an example:

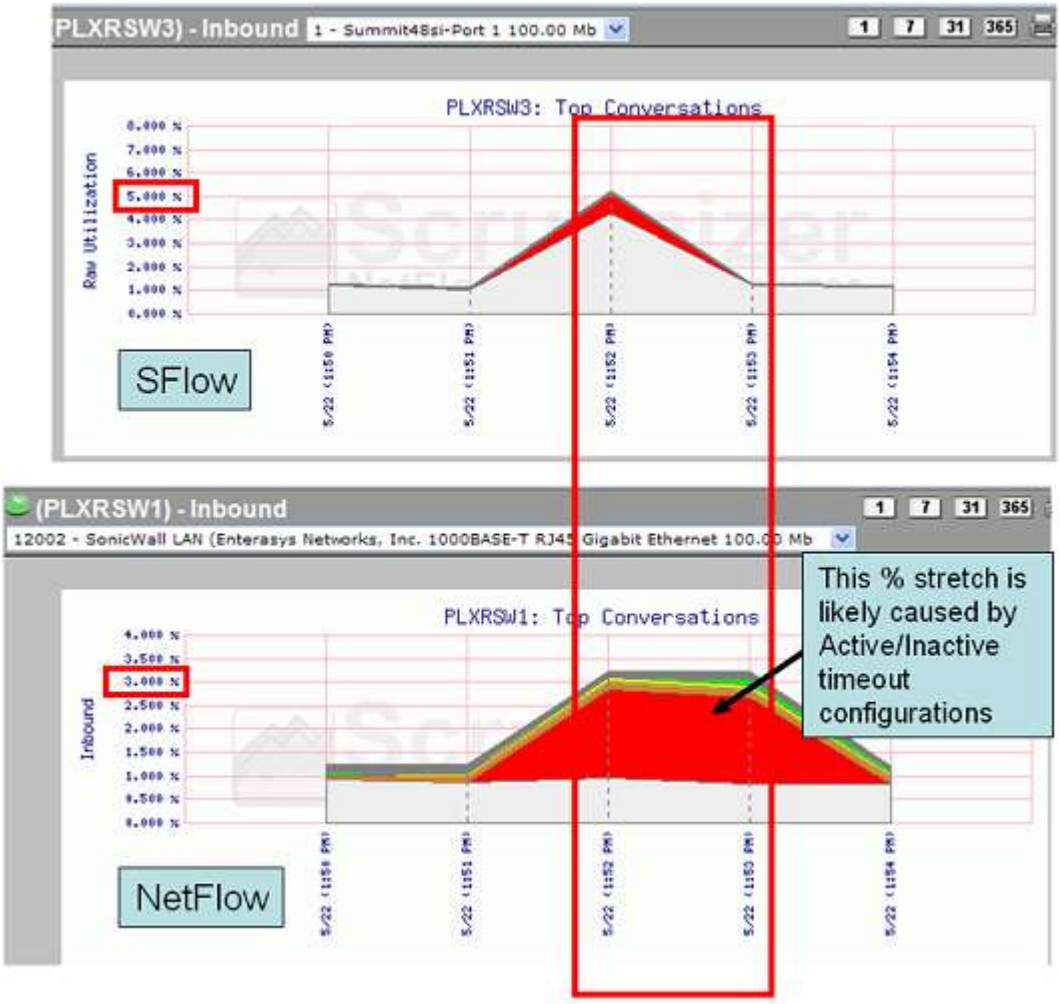

Top Hosts Don't Match Up

The test was left to run for a few days. Scrutinizer sat there collecting away. Every so often the top ten talkers reported were compared for the same time frame. They seldom matched up when looking at trends for the last 5 minutes or the last 24 hours:

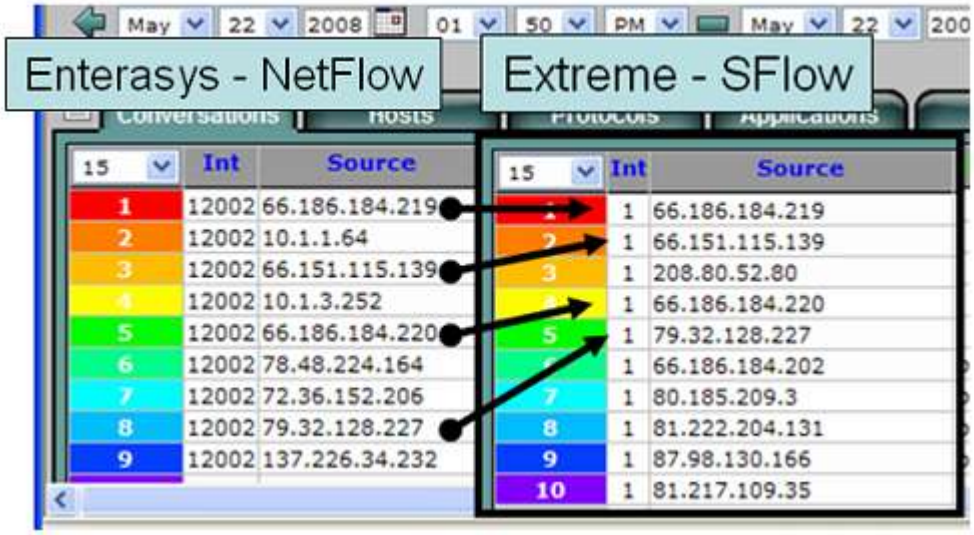

As expected, since the Extreme Summit is sampling packets, the total host bit count is below what the Enterasys Switch is reporting for the same time frame:

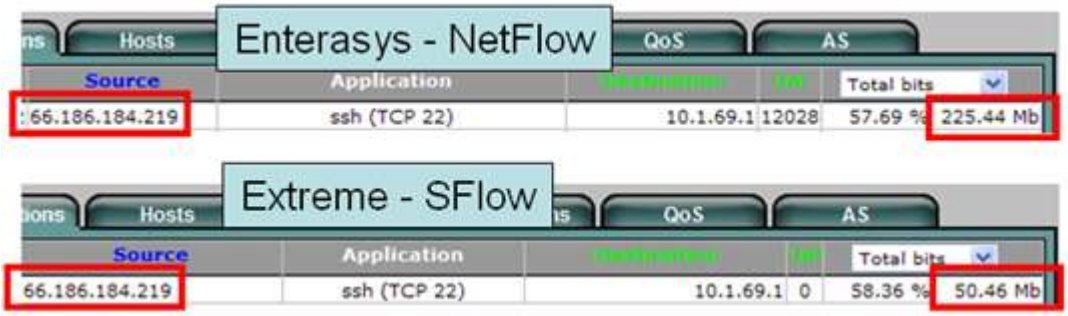

Strictly Speaking IP

When looking at purely IP traffic, NetFlow has the advantage of collecting nearly everything; hence the 4 fold increase over the sFlow interface above. On the other hand, unlike NetFlow, sFlow is not limited to IP traffic and results in more accurate overall utilization. Notice below that the same Outbound traffic reported by NetFlow is under that which is stated by sFlow.

NetFlow Trend:

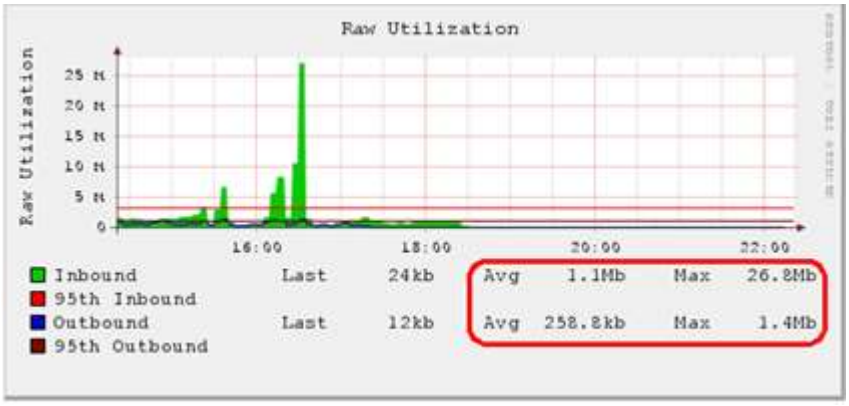

sFlow Trend:

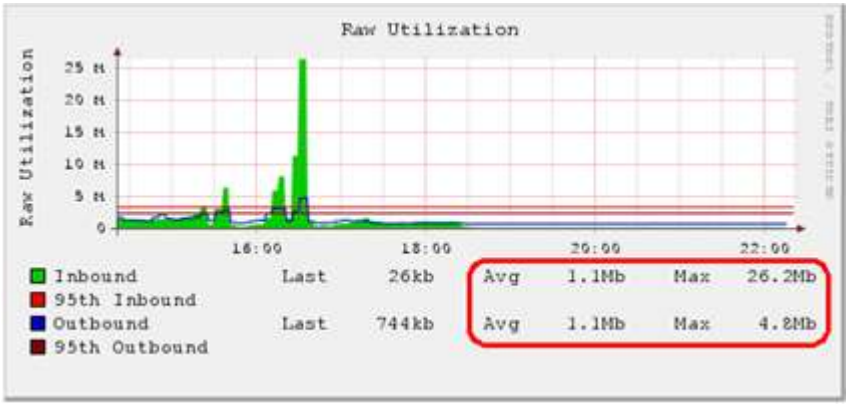

Regarding the above, sFlow reports on non IP traffic as well as broadcasts that are not exported by NetFlow.

"The Enterasys Matrix N-Series switches collect NetFlow statistics for every packet in every flow without sacrificing performance based on the nTERA ASIC capabilities," said Trent Waterhouse - Marketing VP for Enterasys.

"Although we have considered the recent IPFIX solution (based on NetFlow v9), ProCurve currently favors sFlow for unification of our wired and wireless … "

"…. the NetFlow feature is an important transition technology for the "refresh" and we do have plans in our next software release to support NetFlow in our WAN router products." See: Blog with Paul: http://www.networkworld.com/community/node/23982

Flow volumes back to the collector

When the lab technician reviewed the volume of sFlow traffic being sent by the Extreme Summit switch back to the Scrutinizer collector, the results were again interesting. The Extreme sFlow volume was 6 times that of the NetFlow sending Enterasys switch. This is because Plixer configured the Extreme switch to sample as much as possible, which generally isn't necessary. See below:

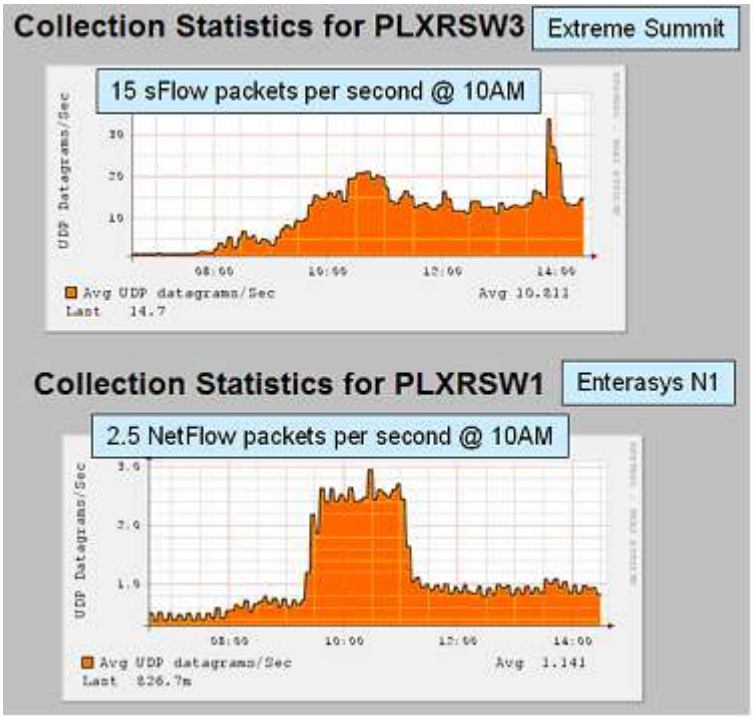

Note that many believe that sFlow is a 1:1 ratio of 1 packet per 1 sample. This is not true. As Wireshark points out below in the packet trace, a single sFlow packet had 8 packet samples in it:

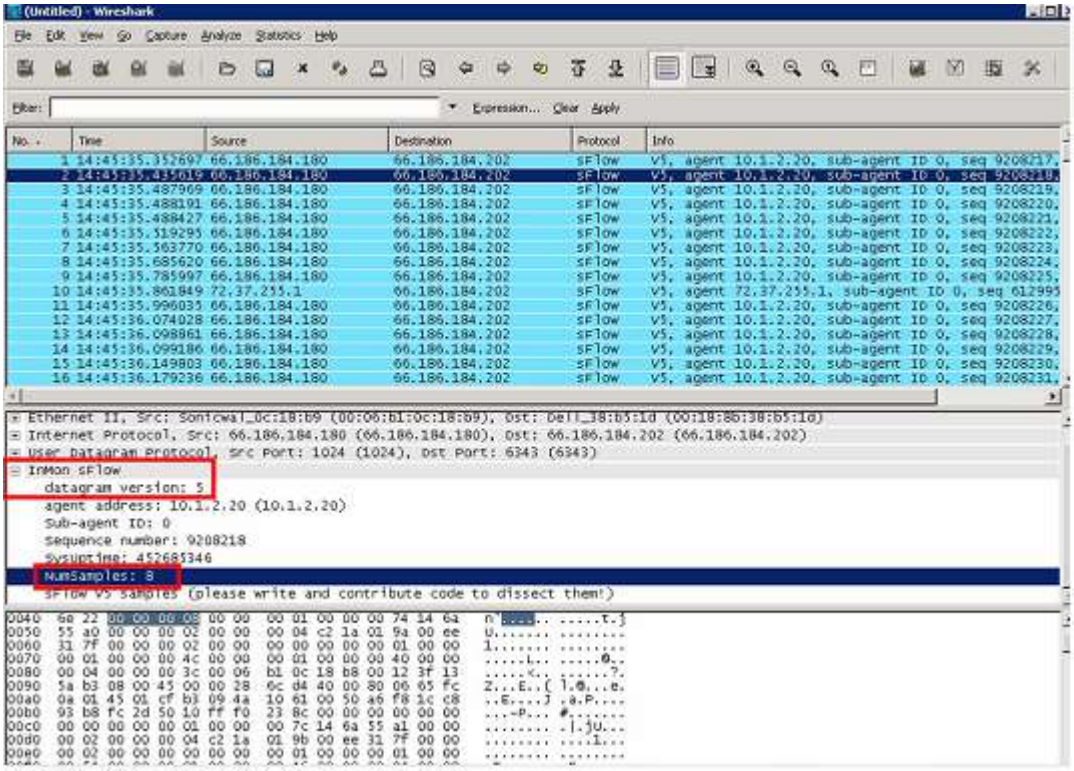

"NetFlow is much more accurate for IP statistics however, sFlow is more than a substitute for NetFlow. It offers many more statistics than NetFlow does. Flexible NetFlow looks to take smart ideas from sFlow like sampling packets."

Marc Bilodeau - CTO, Plixer International, Inc.

Historical Differences

One would think that even with sampling that, statistically, the same top talkers would result with either technology over time and they didn't. Below is based on a 6 day trend on both switches. Although the overall interface utilization trends look the same, the top hosts were inconsistent:

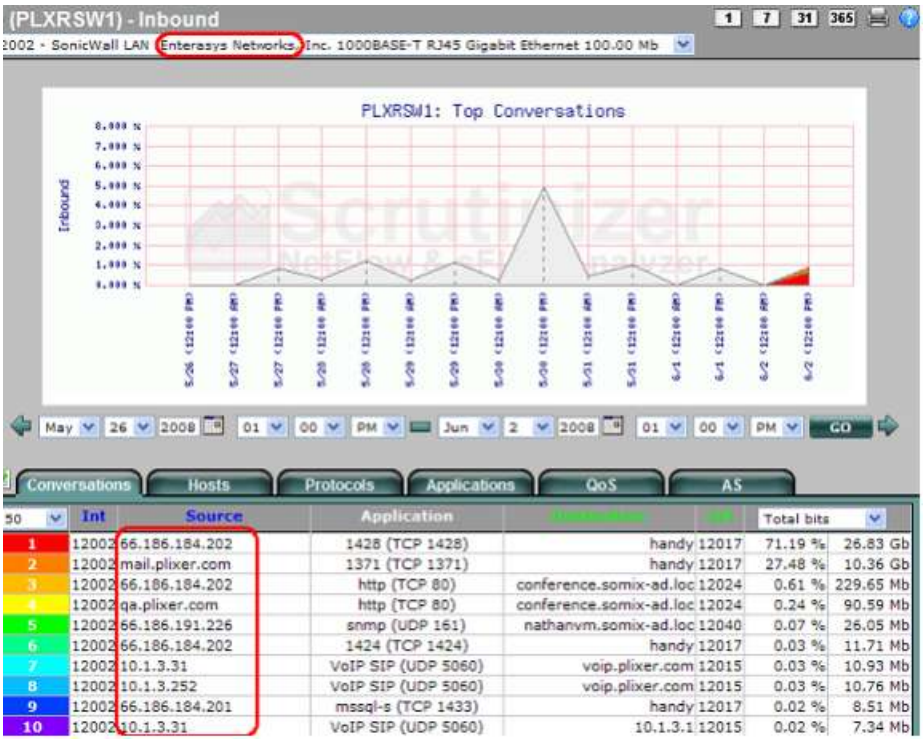

	Int	Source	Application	Destination		Total bits	
1	12002	66.186.184.202	1428 (TCP 1428)	handy	12017	71.19 %	26.83 Gb
2	12002	mail.plixer.com	1371 (TCP 1371)	handy	12017	27.48 %	10.36 Gb
3	12002	66.186.184.202	http (TCP 80)	conference.somix-ad.loc	12024	0.61 %	229.65 Mb
4	12002	qa.plixer.com	http (TCP 80)	conference.somix-ad.loc	12024	0.24 %	90.59 Mb
5	12002	66.186.191.226	snmp (UDP 161)	nathanvm.somix-ad.loc	12040	0.07 %	26.05 Mb
6	12002	66.186.184.202	1424 (TCP 1424)	handy	12017	0.03 %	11.71 Mb
7	12002	10.1.3.31	VoIP SIP (UDP 5060)	voip.plixer.com	12015	0.03 %	10.93 Mb
8	12002	10.1.3.252	VoIP SIP (UDP 5060)	voip.plixer.com	12015	0.03 %	10.76 Mb
9	12002	66.186.184.201	mssql-s (TCP 1433)	handy	12017	0.02 %	8.51 Mb
10	12002	10.1.3.31	VoIP SIP (UDP 5060)	10.1.3.1	12015	0.02 %	7.34 Mb

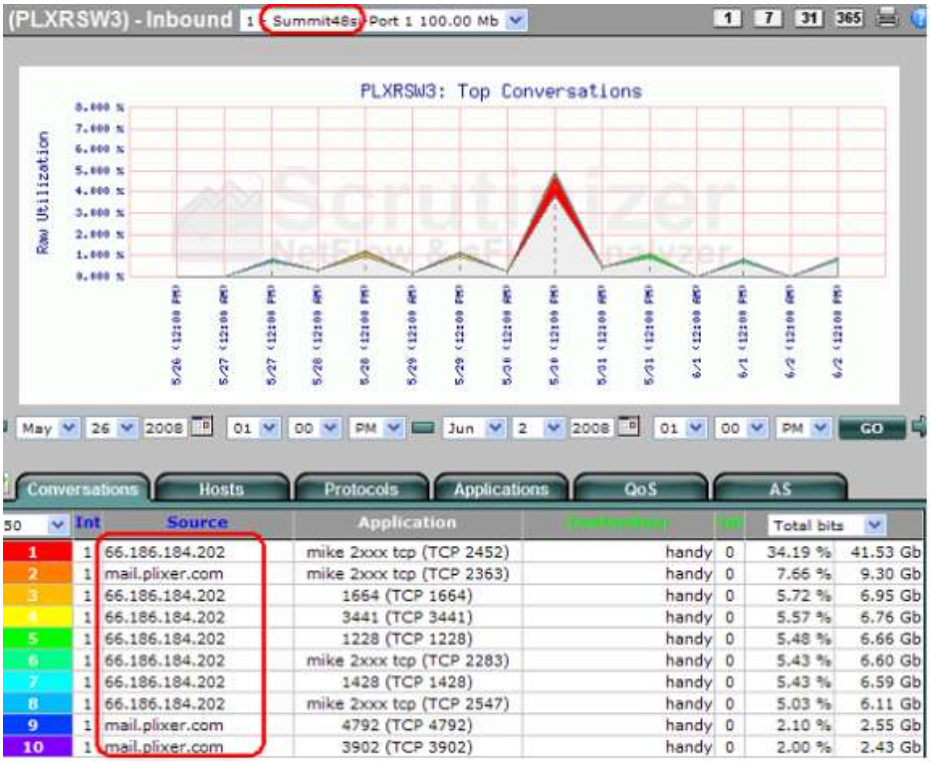

	Int	Source	Application	Destination		Total bits	
1	1	66.186.184.202	mike 2xxx tcp (TCP 2452)	handy	0	34.19 %	41.53 Gb
2	1	mail.plixer.com	mike 2xxx tcp (TCP 2363)	handy	0	7.66 %	9.30 Gb
3	1	66.186.184.202	1664 (TCP 1664)	handy	0	5.72 %	6.95 Gb
4	1	66.186.184.202	3441 (TCP 3441)	handy	0	5.57 %	6.76 Gb
5	1	66.186.184.202	1228 (TCP 1228)	handy	0	5.48 %	6.66 Gb
6	1	66.186.184.202	mike 2xxx tcp (TCP 2283)	handy	0	5.43 %	6.60 Gb
7	1	66.186.184.202	1428 (TCP 1428)	handy	0	5.43 %	6.59 Gb
8	1	66.186.184.202	mike 2xxx tcp (TCP 2547)	handy	0	5.03 %	6.11 Gb
9	1	mail.plixer.com	4792 (TCP 4792)	handy	0	2.10 %	2.55 Gb
10	1	mail.plixer.com	3902 (TCP 3902)	handy	0	2.00 %	2.43 Gb

After comparing the first two switches reporting on the same traffic and seeing inconsistent top 10 host results, plixer decided to review sFlow from a 3rd switch (i.e. the backup plan) looking at the same traffic. The 3rd switch made by Alcatel PLXRSW2 was sampling at a much lower rate, but the top ten hosts were consistent with the Extreme sFlow switch:

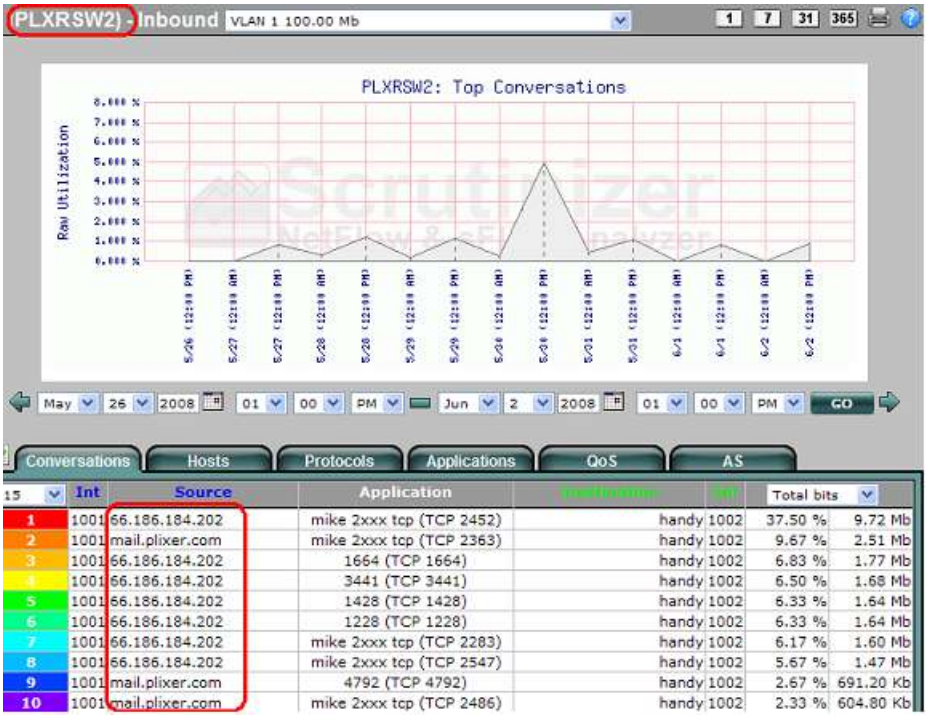

Conclusion

Both technologies have their benefits. Because of the cost involved with engineering NetFlow on a switch and the readily available sFlow chips from Inmon, sFlow is the prevailing technology on switches. On routers, NetFlow seems to be the more popular technology.

In extremely high traffic volume environments, sampling is the only alternative as no collector can handle the volume of flows generated by even a single router. Even Cisco recommends sampling albeit with NetFlow v9.

Related Articles:

Cisco toe stepper HP ProCurve deftly hoofs over Cisco NetFlow
http://www.networkworld.com/community/node/23982

Cisco's NetFlow vs. Inmon's sFlow: Which will prevail?
http://www.networkworld.com/community/node/22667

NetFlow or sFlow: which is the open standard?
http://www.networkworld.com/community/node/23739