Plixer

WHITE PAPER

NetFlow vs. sFlow: a technical review

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 used. 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, but an architecture question of where should each be deployed. NetFlow (i.e. IPFIX) 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. Many vendors support NetFlow.

sFlow

Sflow is a packet-sampling technology. Some sFlow 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, it 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: sflow.org/products/network.php

Developments in NetFlow v9 allow it to sample similarly 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.

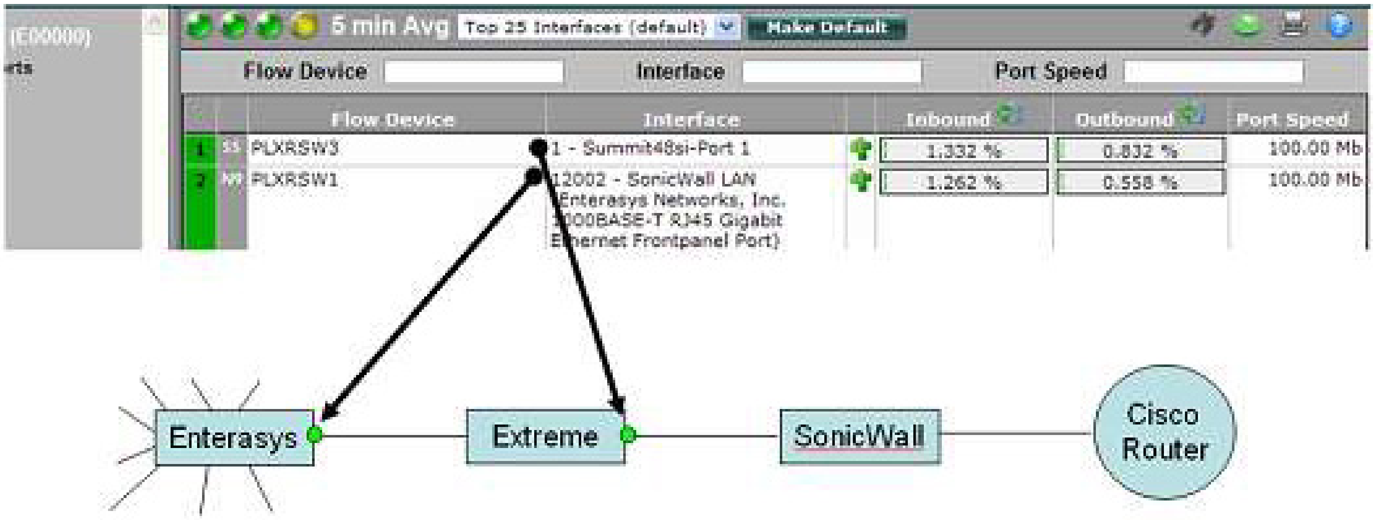

(E00000)

rts

5 min Avg | Top 25 Interfaces (default) | Make Default

Flow Device | Interface | Port Speed

		Flow Device	Interface		Inbound	Outbound	Port Speed
1	SS	PLXRSW3	1 - Summit48si-Port 1		1.332 %	0.832 %	100.00 Mb
2	NX	PLXRSW1	12002 - SonicWall LAN (Enterasys Networks, Inc. 1000BASE-T RJ45 Gigabit Ethernet Frontpanel Port)		1.262 %	0.558 %	100.00 Mb

Collection and analysis

For flow collection, Scrutinizer NetFlow & sFlow Analyzer v6 was used, which is pictured above. 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 likely had many factors, 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. As a result, 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 nearly the same.

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:

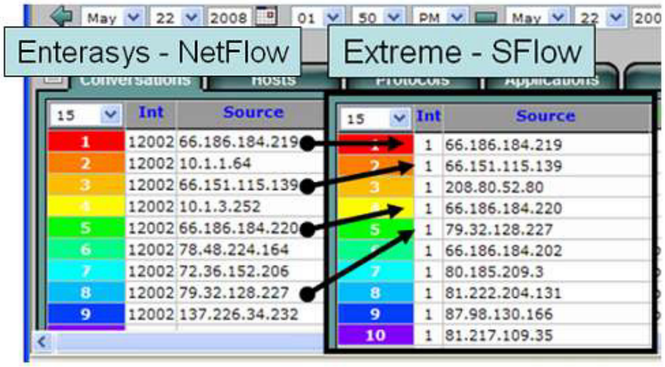

May 22 2008 01 50 PM — May 22 200

Enterasys - NetFlow

15	Int	Source
1	12002	66.186.184.219
2	12002	10.1.1.64
3	12002	66.151.115.139
4	12002	10.1.3.252
5	12002	66.186.184.220
6	12002	78.48.224.164
7	12002	72.36.152.206
8	12002	79.32.128.227
9	12002	137.226.34.232

Extreme - SFlow

15	Int	Source
1	1	66.186.184.219
2	1	66.151.115.139
3	1	208.80.52.80
4	1	66.186.184.220
5	1	79.32.128.227
6	1	66.186.184.202
7	1	80.185.209.3
8	1	81.222.204.131
9	1	87.98.130.166
10	1	81.217.109.35

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 host for the same timeframe:

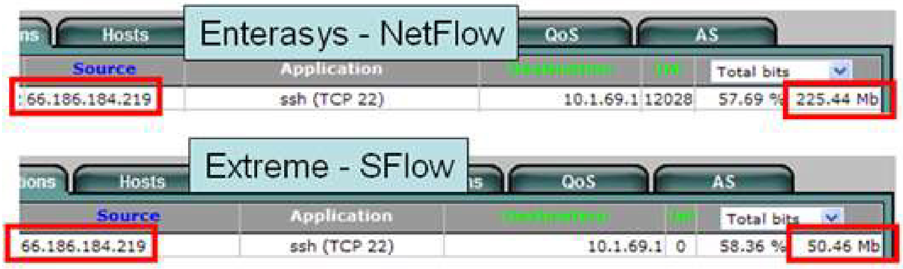

Enterasys - NetFlow

Source	Application	Destination	Int	Total bits	
66.186.184.219	ssh (TCP 22)	10.1.69.1	12028	57.69 %	225.44 Mb

Extreme - SFlow

Source	Application	Destination	Int	Total bits	
66.186.184.219	ssh (TCP 22)	10.1.69.1	0	58.36 %	50.46 Mb

Strictly speaking IP

When looking purely at IP traffic, NetFlow has the advantage of collecting nearly everything; hence the fourfold 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 outbound traffic reported by NetFlow is lower than that stated by sFlow.

NetFlow Trend:

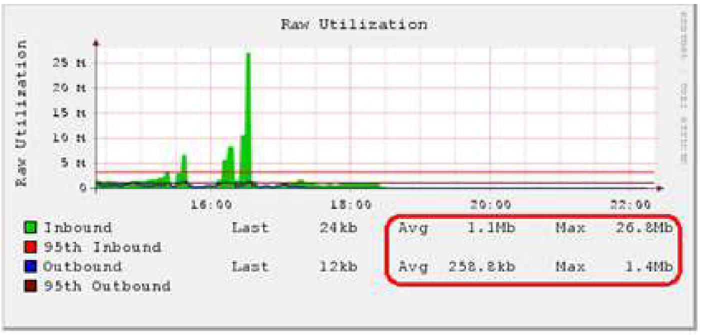

sFlow Trend:

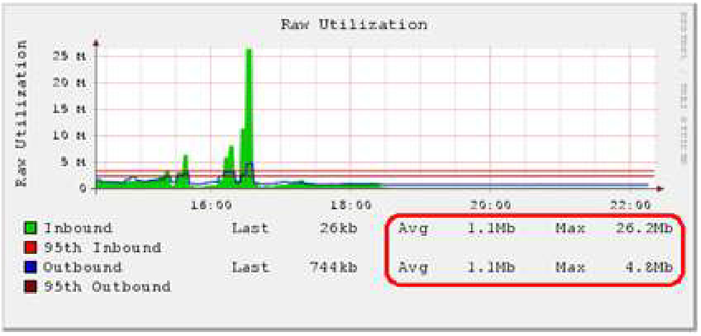

Regarding the above, sFlow reports on non-IP traffic, as well as broadcasts that are not exported by NetFlow.

"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." Source

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 six times that of the NetFlow-sending Enterasys switch. This is because Plixer configured the Extreme switch to sample as much as possible, which usually isn't necessary. See below:

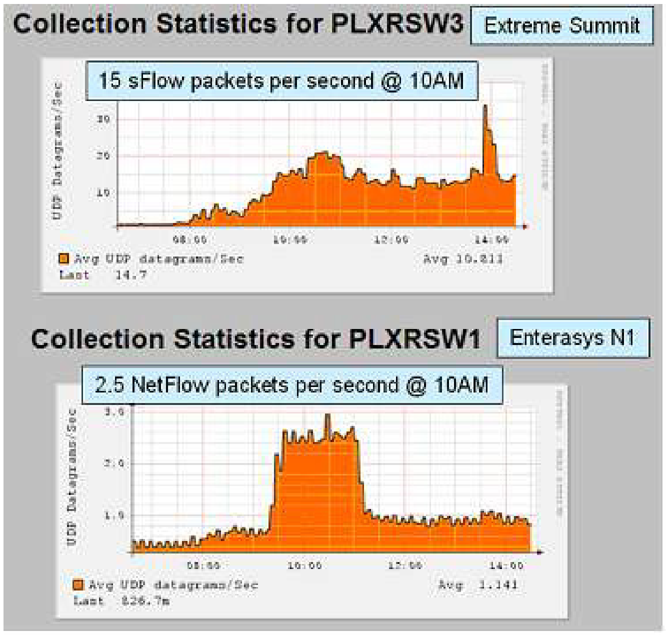

Note that many believe that sFlow is a 1:1 ratio of 1 packet per 1 sample. This is not true. As Wireshark points out in the packet trace on the following page, a single sFlow packet had 8 packet samples in it.

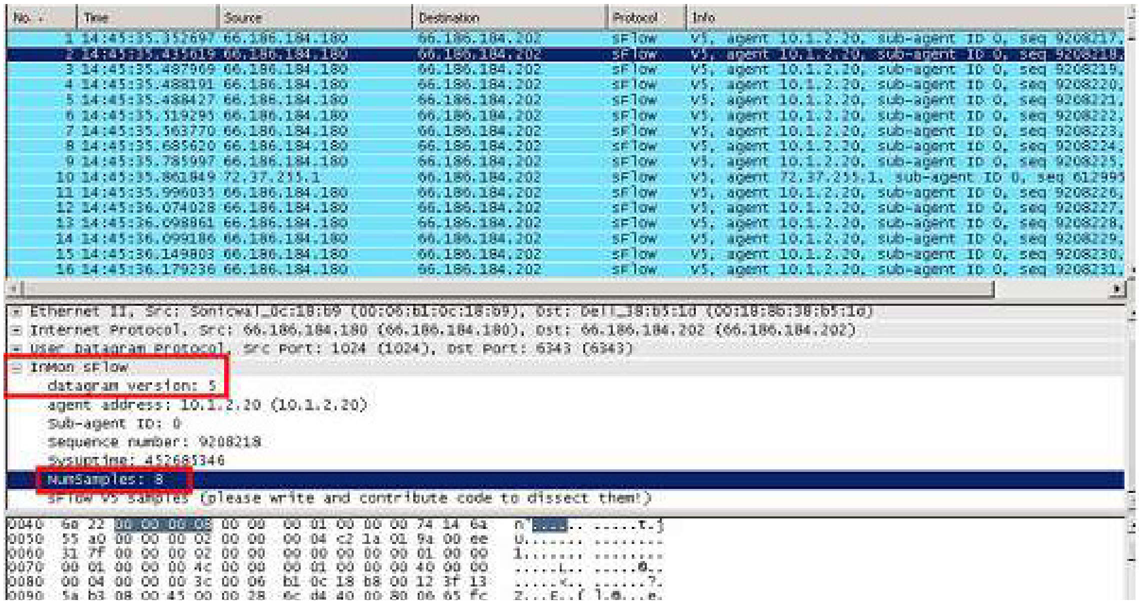

```
No. .   Time             Source           Destination      Protocol  Info
      1 14:45:35.352697 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208217,
      2 14:45:35.435619 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208218,
      3 14:45:35.487969 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208219,
      4 14:45:35.488191 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208220,
      5 14:45:35.488427 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208221,
      6 14:45:35.519295 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208222,
      7 14:45:35.563770 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208223,
      8 14:45:35.685620 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208224,
      9 14:45:35.785997 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208225,
     10 14:45:35.861849 72.37.235.1      66.186.184.202   sFlow     V5, agent 72.37.235.1, sub-agent ID 0, seq 612995
     11 14:45:35.996035 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208226,
     12 14:45:36.074028 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208227,
     13 14:45:36.098861 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208228,
     14 14:45:36.099186 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208229,
     15 14:45:36.149803 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208230,
     16 14:45:36.179236 66.186.184.180   66.186.184.202   sFlow     V5, agent 10.1.2.20, sub-agent ID 0, seq 9208231,

⊞ Ethernet II, Src: Sonicwal_0c:18:b9 (00:06:b1:0c:18:b9), Dst: Dell_38:b5:1d (00:18:8b:38:b5:1d)
⊞ Internet Protocol, Src: 66.186.184.180 (66.186.184.180), Dst: 66.186.184.202 (66.186.184.202)
⊞ User Datagram Protocol, Src Port: 1024 (1024), Dst Port: 6343 (6343)
⊟ InMon sFlow
    datagram version: 5
    agent address: 10.1.2.20 (10.1.2.20)
    Sub-agent ID: 0
    Sequence number: 9208218
    SysUptime: 452685346
    NumSamples: 8
    sFlow V5 samples (please write and contribute code to dissect them!)

0040  6a 22 00 00 00 08 00 00  00 01 00 00 00 74 14 6a   n"...... .....t.j
0050  55 a0 00 00 00 02 00 00  00 04 c2 1a 01 9a 00 ee   U....... ........
0060  31 7f 00 00 00 02 00 00  00 00 00 00 00 01 00 00   1....... ........
0070  00 01 00 00 00 4c 00 00  00 01 00 00 00 40 00 00   .....L.. .....@..
0080  00 04 00 00 00 3c 00 06  b1 0c 18 b8 00 12 3f 13   .....<.. ......?.
0090  5a b3 08 00 45 00 00 28  6c d4 40 00 80 06 65 fc   Z...E..( l.@...e.
```

> "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

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:

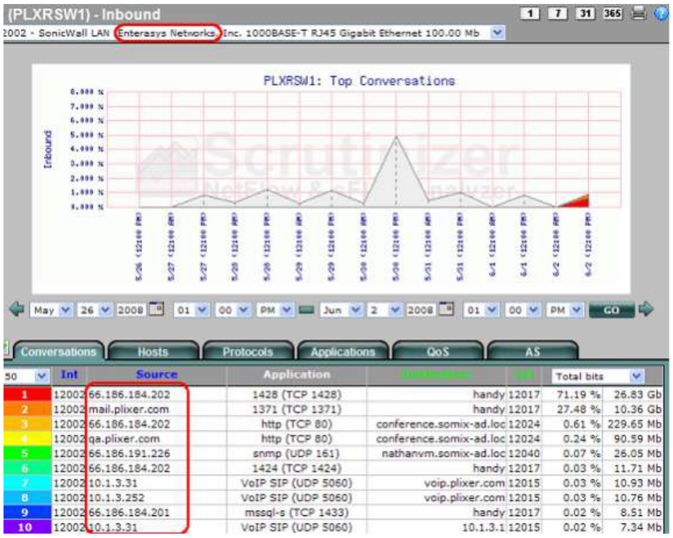

(PLXRSW1) - Inbound

2002 - SonicWall LAN Enterasys Networks, Inc. 1000BASE-T RJ45 Gigabit Ethernet 100.00 Mb

50	Int	Source	Application	Destination	If	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

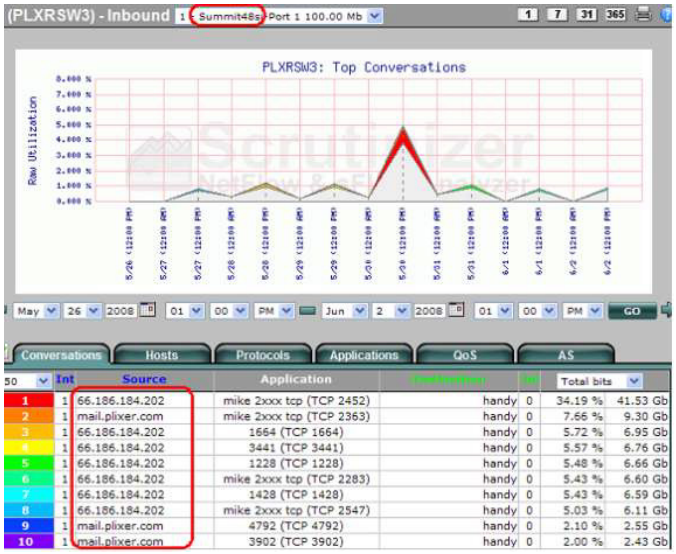

(PLXRSW3) - Inbound

1 - Summit48si Port 1 100.00 Mb

50	Int	Source	Application	Destination	If	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 third switch (i.e. the backup plan) looking at the same traffic.

The third switch, PLXRSW2, made by Alcatel, was sampling at a much lower rate, but the top ten hosts were consistent with the Extreme sFlow switch.

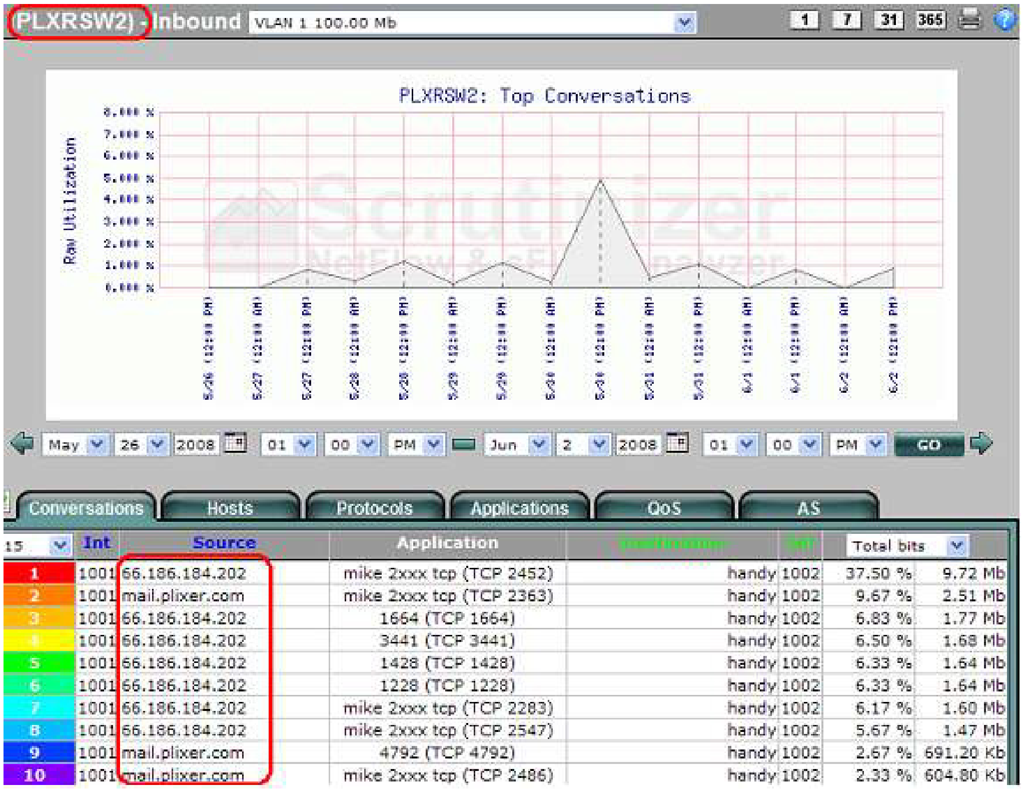

15	Int	Source	Application	Destination	Int	Total bits	
1	1001	66.186.184.202	mike 2xxx tcp (TCP 2452)	handy	1002	37.50 %	9.72 Mb
2	1001	mail.plixer.com	mike 2xxx tcp (TCP 2363)	handy	1002	9.67 %	2.51 Mb
3	1001	66.186.184.202	1664 (TCP 1664)	handy	1002	6.83 %	1.77 Mb
4	1001	66.186.184.202	3441 (TCP 3441)	handy	1002	6.50 %	1.68 Mb
5	1001	66.186.184.202	1428 (TCP 1428)	handy	1002	6.33 %	1.64 Mb
6	1001	66.186.184.202	1228 (TCP 1228)	handy	1002	6.33 %	1.64 Mb
7	1001	66.186.184.202	mike 2xxx tcp (TCP 2283)	handy	1002	6.17 %	1.60 Mb
8	1001	66.186.184.202	mike 2xxx tcp (TCP 2547)	handy	1002	5.67 %	1.47 Mb
9	1001	mail.plixer.com	4792 (TCP 4792)	handy	1002	2.67 %	691.20 Kb
10	1001	mail.plixer.com	mike 2xxx tcp (TCP 2486)	handy	1002	2.33 %	604.80 Kb

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
networkworld.com/community/node/23982

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

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

The content in this document is for informational purposes only and is subject to change by Plixer, LLC without notice. While reasonable efforts have been made in the preparation of this publication to ensure its accuracy, Plixer, LLC assumes no liability resulting from technical or editorial errors or omissions, or for any damages resulting from the use of this information. Unless specifically included in a written agreement with Plixer, Plixer, LLC has no obligation to develop or deliver any future release or upgrade or any feature, enhancement, or function.