WHITEPAPER

Best Practices for NetFlow/IPFIX Analysis and Reporting

IT managers and network administrators are constantly making decisions affecting critical business activity on the network. Management challenges include solving performance problems faster, optimizing the network infrastructure for application performance, proving the performance of applications delivered over the enterprise network, making more informed infrastructure investments, and many others.

The key to confidently managing network issues is complete visibility into the network – the ability to measure, quantify and analyze enterprise traffic across the network. In the past, network information has been gathered using hardware probes that collected RMON2 statistics. Today, with Cisco IOS® NetFlow, and the emerging IPFIX standard, existing routers and switches within network infrastructures can continuously collect flow data.

This paper examines cost-effective NetFlow/IPFIX reporting and analysis solutions, and provides examples of how to use this data to help better manage the network for application performance.

Importance of Traffic Analysis

Network engineers that make decisions about network changes do it best when they understand the traffic behavior on their network. Network data and some of the analyses that make this possible are:

- » Real-time traffic volume and rate by application protocol, host, and host-to-host (conversation). This data can be collected over time and refined to produce even more meaningful information, such as a profile of typical network operation.
- » Traffic sources that are grouped by business unit, geography, subnets, or other helpful way. The network administrators can then associate network traffic with business entities or functions, and trend growth per grouping.
- » Filtering traffic in order to see anomalies. For instance, an average traffic rate over a 24-hour period might not allow the spike that occurs during working hours to be noticed for some time. Unused traffic hours during the nighttime must be filtered to produce the timeframe of interest when most users are active. This filtering is important when setting threshold alarms.

Further analysis can lead to capacity planning, network readiness assessments for new business application rollouts, accounting for network usage, identifying and eliminating unwanted traffic, planning and implementing QoS, and performing budget assessments for future network equipment and support. Without comprehensive knowledge of the traffic, network, and application performance can be unpredictable, unacceptable, and costly. Network administration and engineering changes become reactive, only being put into service quickly when something goes wrong. Projected costs for rolling out new applications or adding new sites may not be accurately estimated. Possible overrun on IT budgets might well be avoided with usage-based planning.

Network Instrumentation

Instrumenting the network involves identifying strategic points where data could most profitably be gathered. In the past, data-gathering probes were placed at remote sites, at points of Layer 2 or Layer 3 aggregation, or both using a hybrid model. This common approach to instrument and measure traffic on the network using RMON2 probes was used extensively. Today, with NetFlow/IPFIX available on a wide-range of routers and switches, IT managers and CIO's have a more cost-effective alternative.

RMON2 Network Probes

RMON2 probes are dedicated instruments that monitor data packets as they cross the network at certain key aggregation points. By observing and measuring RMON2 data packet behavior on the network, network probes gather protocol and application performance-related data. The main advantages of deploying probes were:

- » Real-time traffic data captured
- » No additional load introduced on the network
- » Information gathered on a wide range of protocols, such as TCP/IP, UDP/IP, ICMP, IPX, and NetBEUI

Probe implementations typically have a central management station where data is aggregated and reported across all probes in the network. Such aggregations provide granular, as well as global, views of network traffic. RMON2 probes can provide rich network performance information; however, the following characteristics are also associated with their deployment:

- High capital investment˜ how many probes will be needed to cover all or part of the network and how much will that level of coverage cost?
- Ongoing labor costṣ˜ to perform configuration, planning, and deployment of probes
- Disruptive installation costs – probe installation requires breaking the network circuit to Install.
- Non-scalable setup costs – most probe solutions require that each probe be configured for each monitored traffic type.
- Medium to high lifecycle maintenance costs – licensing, software upgrades, probe interface upgrades as network bandwidth increases.

For parts of a network that are already instrumented with probes, an extensive polling and reporting solution for network management may be required to leverage existing hardware. For future investment, consider the NetFlow/IPFIX alternative described below.

NetFlow/IPFIX

With the increasing popularity of Cisco's NetFlow technology, a standard is emerging for similarly exporting traffic flows. Due for final approval this year, the Internet Engineering Task Force (IETF) is standardizing flow export under RFC 3917 – IP Flow Information Export (IPFIX). Using Cisco IOS NetFlow v9 as a foundation, the IPFIX standard allows vendors to consistently provide flow data exports in the same format. With this technology, existing routers are used to gather NetFlow/IPFIX data, providing enterprises with the following advantages:

- No capital investment – almost all networks already contain devices with NetFlow/IPFIX capability. By simply turning on the NetFlow/IPFIX data, these devices can immediately begin exporting network statistics.
- Low deployment costs – configuring NetFlow/IPFIX involves a few global commands and an interface command for each interface running NetFlow/IPFIX.
- Complete data source – NetFlow/IPFIX measures and reports automatically on all IP traffic.
- No lifecycle maintenance – NetFlow/IPFIX is associated with Cisco router hardware/software maintenance.

When looking to deploy NetFlow/IPFIX, consider:

- NetFlow/IPFIX typically increases CPU utilization on the configured devices on average by only 1% to 2%.
- Only IP traffic is supported.

Using NetFlow/IPFIX Data

Each NetFlow/IPFIX flow is unique and is identified by seven criteria: Source IP address, Destination IP address, Source Port number (TCP/UDP), Destination Port number (TCP/UDP), Layer 3 Protocol Type (IP/ICMP),Type of Service (ToS), and Input logical interface. Any variation in these criteria distinguishes one flow from another.
NetFlow/IPFIX data can be analyzed to report:

» All hosts that transmit the most data on the network
» All hosts that transmit the most data to each other
» All applications that put the most traffic on the network
» Data volumes per entity (circuit, remote location, region, etc.)
» Data rates per entity (circuit, remote location, region, etc.)
» ToS markings per application or entity

NetFlow/IPFIX Solution Criteria

Reporting on NetFlow/IPFIX data requires a robust analysis engine. When selecting your NetFlow/IPFIX reporting solution, consider the following criteria.

Scalability

A NetFlow/IPFIX management/reporter solution should be able to scale from small implementations such as a data center and a few remote offices to large enterprise implementations with multiple data centers and thousands of remote sites. As with routing protocols, a hierarchical solution provides the most scalability. In a hierarchical solution, NetFlow/IPFIX data originates at the router. Sending NetFlow/IPFIX data to a common collection device aggregates router information together. Each collection device reduces the NetFlow/IPFIX data and forwards it to the main management/ reporting console. This has the dual effect of minimizing the data flowing across the network and maximizing the network area that can be covered with a NetFlow/IPFIX solution.

At the same time, each collection device must be careful not to discard too much information, rendering the solution useless or ineffective at gaining visibility into the network. Ideally, a complete reporting solution collects thousands of unique protocols, hosts, and conversations per network link in addition to 100% of data for mission-critical applications. Optimal NetFlow implementations also allow network managers to minimize the number of data harvesters and decrease costs while maximizing the coverage of the network.

Broad Reporting Capabilities

The NetFlow management/reporting station should not only allow enterprise-level visibility, but provide application specific details when requested. For this depth, the solution should provide an in-depth analysis of NetFlow/IPFIX data from granular information such as reports by particular hosts and protocols to aggregated information such as reports by transport or application protocol, interface aggregations, circuit, geographical region, or business entity. The reporting solution should also give the user the ability to run complex queries on historical data, provide troubleshooting insight with real-time data, and allow further analysis on every raw NetFlow datagram in the network. The NetFlow/IPFIX management/reporting station must be able to quickly produce reports and respond to these queries.

The NetQoS NetFlow/IPFIX Solution

NetQoS, Inc. has developed the standard for enterprise NetFlow/IPFIX reporting called NetQoS ReporterAnalyzer. Working closely with Cisco, this scalable solution has been developed to take advantage of the rich set of data that NetFlow/IPFIX offers. Several hardware/software components comprise the system. NetFlow Harvesters are placed in close proximity to devices to collect and process NetFlow datagrams, NetFlow/IPFIX Manager aggregates and translates NetFlow/IPFIX statistics to the Data Storage Appliances, and ReporterAnalyzer presents the information in easy-to-read web based reports. This combination of components can monitor and report on activity from over 100,000 interfaces. With this kind of scalability, ReporterAnalyzer helps eight of the Fortune 12, and twenty-four of the Fortune 50 enterprises actively manage the world's most demanding networks.

Solution capabilities include:

» High-level Operations View – Top interfaces, protocols, and hosts from across your enterprise displayed on a single page. Quickly determine the top "hot spots" on your network.

» Revealing Traffic Analysis Reports – Rate, volume, and utilization measurements by protocol, host, and conversation for a full year. Everything you need to know about who, what, where, when, and how much.

» Application information defined by a combination of ports, IP addresses, and ToS. Determine the amount of VoIP or other particular traffic, and who is generating it.

» Real-time reports and alarms at one-minute granularity for every interface. Know when predefined threshold violations occur at the time they occur. Identify worms, viruses, and denial of service attacks and the hosts they infect.

» Set rate, volume, and count thresholds based on flows, packets, or bytes.

» Track 100% of all flow traffic on your network. Pinpoint analysis of any network occurrence.

» View custom reports by grouping interface, protocols, hosts, and conversations. Particularly important when drilling down quickly into broad issues.

» Set thresholds based on past performance. Confidently and proactively plan for network upgrades when the time approaches, before problems occur.

Samples of NetFlow/IPFIX Data Analysis using ReporterAnalyzer

The following three sample scenarios were gathered from NetQoS customers. They demonstrate how NetFlow/IPFIX data supports the extensive reporting capability and in-depth analysis of ReporterAnalyzer to empower the network engineer in making informed decisions.

Example 1: Making Informed Infrastructure Investments: What is the source of the performance problem in Singapore? Will a $120,000 upgrade in bandwidth fix the problem?

Problem: Users in Singapore are complaining about poor performance when accessing critical business applications in the Houston Data Center. The network manager has suggested a $120,000 bandwidth upgrade to fix the problem. However, the network engineer is not confident that this will resolve the issue. The network engineer uses ReporterAnalyzer to understand the cause of the poor performance.

Analysis: The network engineer views the Enterprise Overview page for Interface Utilization. See **Figure 1**.

Figure 1

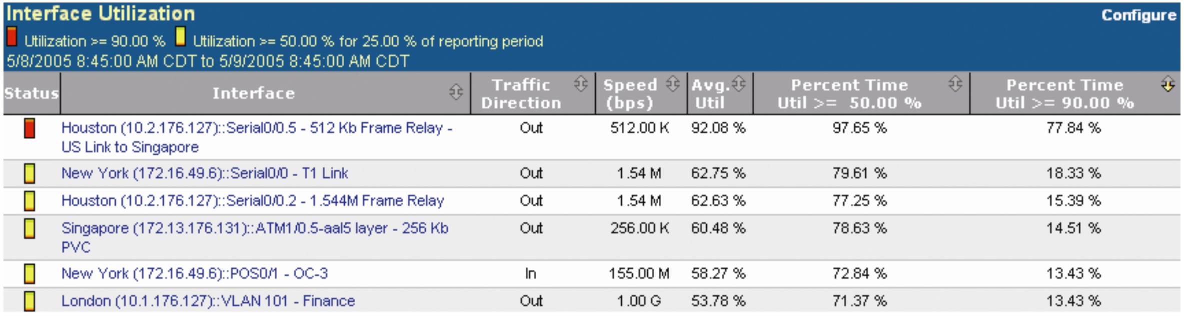

Interface Utilization

Configure

🟥 Utilization >= 90.00 % 🟨 Utilization >= 50.00 % for 25.00 % of reporting period

5/8/2005 8:45:00 AM CDT to 5/9/2005 8:45:00 AM CDT

Status	Interface	Traffic Direction	Speed (bps)	Avg. Util	Percent Time Util >= 50.00 %	Percent Time Util >= 90.00 %
🟥	Houston (10.2.176.127)::Serial0/0.5 - 512 Kb Frame Relay - US Link to Singapore	Out	512.00 K	92.08 %	97.65 %	77.84 %
🟨	New York (172.16.49.6)::Serial0/0 - T1 Link	Out	1.54 M	62.75 %	79.61 %	18.33 %
🟨	Houston (10.2.176.127)::Serial0/0.2 - 1.544M Frame Relay	Out	1.54 M	62.63 %	77.25 %	15.39 %
🟨	Singapore (172.13.176.131)::ATM1/0.5-aal5 layer - 256 Kb PVC	Out	256.00 K	60.48 %	78.63 %	14.51 %
🟨	New York (172.16.49.6)::POS0/1 - OC-3	In	155.00 M	58.27 %	72.84 %	13.43 %
🟨	London (10.1.176.127)::VLAN 101 - Finance	Out	1.00 G	53.78 %	71.37 %	13.43 %

The US to Singapore link is showing a red status, indicating its utilization is over 90% in the interface utilization table and the Top Out Interfaces bar chart. See **Figure 2**. The engineer clicks on the link name to see more information.

Figure 2

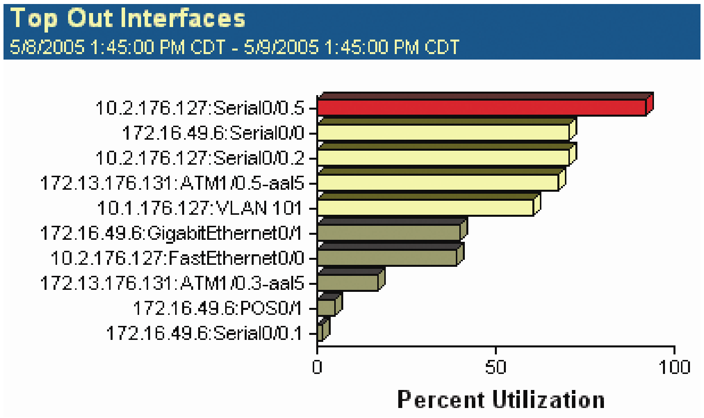

The engineer gets curious about how long this interface has been above 90% utilization. He clicks into the interface and displays a Calendar Chart for the last 30 days on that interface. The display is shown in **Figure 3**.

Figure 3

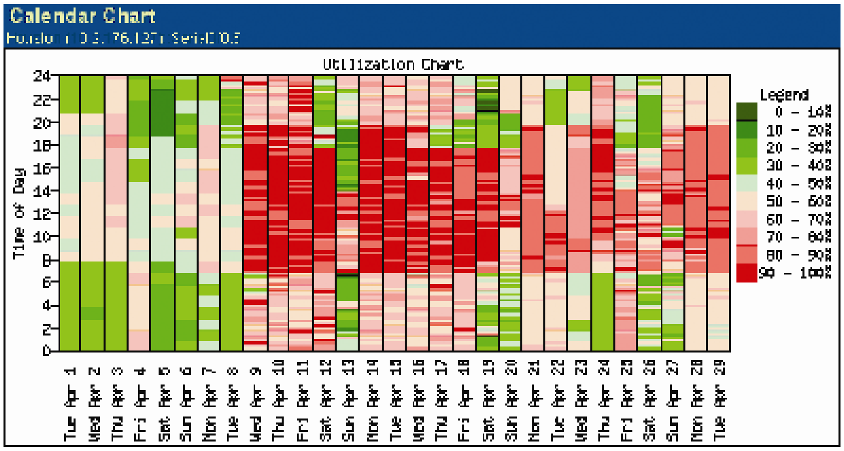

The red color indicates greater than 90% utilization; it began just before midnight on April 8th. The network engineer then wants an understanding of what type of traffic is causing the high utilization. So he accesses a Top Applications chart for that interface, as shown in **Figure 4**.

Figure 4

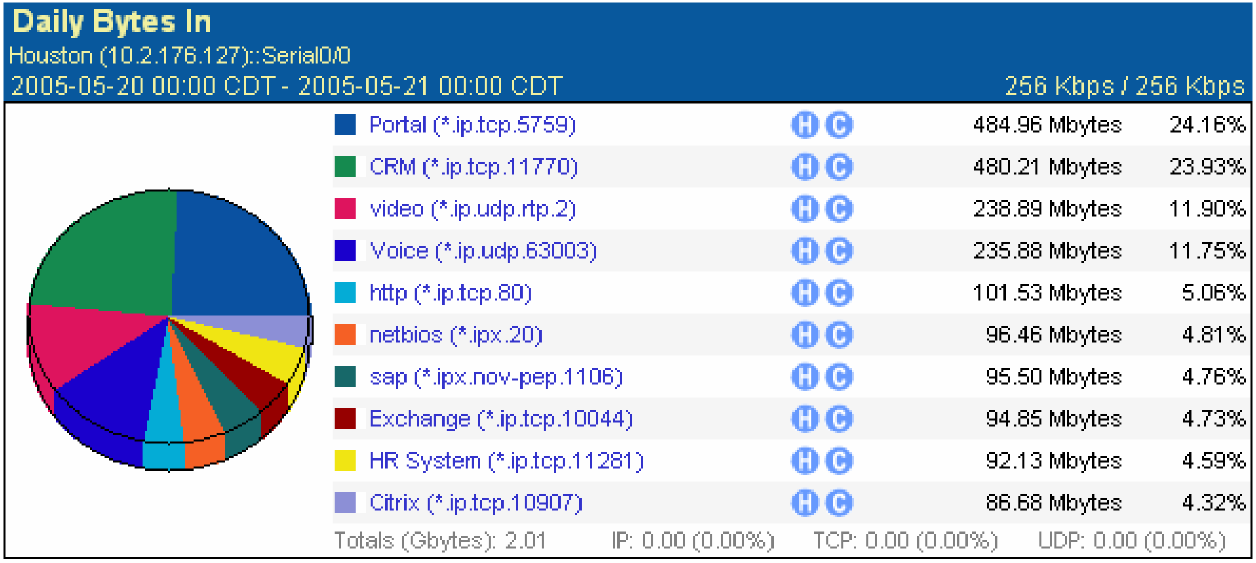

Daily Bytes In

Houston (10.2.176.127)::Serial0/0

2005-05-20 00:00 CDT - 2005-05-21 00:00 CDT 256 Kbps / 256 Kbps

🟦 Portal (*.ip.tcp.5759)	H C	484.96 Mbytes	24.16%
🟩 CRM (*.ip.tcp.11770)	H C	480.21 Mbytes	23.93%
🟥 video (*.ip.udp.rtp.2)	H C	238.89 Mbytes	11.90%
🟦 Voice (*.ip.udp.63003)	H C	235.88 Mbytes	11.75%
🟦 http (*.ip.tcp.80)	H C	101.53 Mbytes	5.06%
🟧 netbios (*.ipx.20)	H C	96.46 Mbytes	4.81%
🟩 sap (*.ipx.nov-pep.1106)	H C	95.50 Mbytes	4.76%
🟥 Exchange (*.ip.tcp.10044)	H C	94.85 Mbytes	4.73%
🟨 HR System (*.ip.tcp.11281)	H C	92.13 Mbytes	4.59%
🟪 Citrix (*.ip.tcp.10907)	H C	86.68 Mbytes	4.32%

Totals (Gbytes): 2.01 IP: 0.00 (0.00%) TCP: 0.00 (0.00%) UDP: 0.00 (0.00%)

He sees that http is by far the largest component of the protocols going out on the interface. The engineer clicks on the H in the blue circle to the right of the protocol name. This opens a chart of the hosts using this protocol. See **Figure 5**.

Figure 5

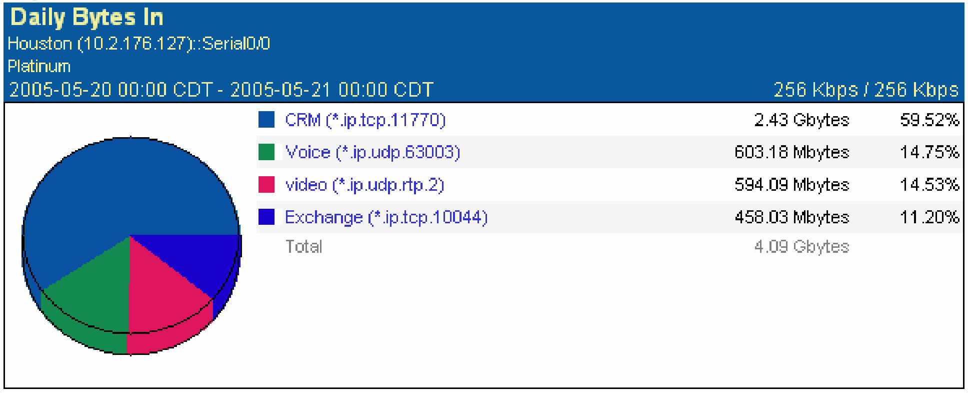

Daily Bytes In
Houston (10.2.176.127)::Serial0/0
Platinum
2005-05-20 00:00 CDT - 2005-05-21 00:00 CDT 256 Kbps / 256 Kbps

🟦 CRM (*.ip.tcp.11770)	2.43 Gbytes	59.52%
🟩 Voice (*.ip.udp.63003)	603.18 Mbytes	14.75%
🟥 video (*.ip.udp.rtp.2)	594.09 Mbytes	14.53%
🟦 Exchange (*.ip.tcp.10044)	458.03 Mbytes	11.20%
Total	4.09 Gbytes	

He sees that the US Web Proxy Server is handling almost 83% of the http traffic. He has now found the source of the problem. The users in Singapore have changed their Internet proxy server to the US Web Proxy server from the local Singapore proxy server to get faster Internet access. The engineer tells the local support group in Singapore to change the proxy server settings back to the local proxy server.

Solution: In this example global visibility with ReporterAnalyzer helped the network engineer troubleshoot how misconfigured Internet traffic was impacting network performance. In less than five minutes with ReporterAnalyzer, the network engineer was able to confidently tell the IT Manager that the costly upgrade was unnecessary. Using ReporterAnalyzer spared the company a $120,000 expense.

Example 2: Solving Problems Faster: Why can't users gain access to key financial applications or server resources?

Problem: In the last ten minutes, users in the New York office have suddenly started calling to complain that they can't access key financial applications or server resources in the London data center.

Analysis: The network engineer receiving the calls needs to know what is on the network right now. The engineer opens a real-time report on the appropriate network interface, as shown in Figure 6.

Figure 6

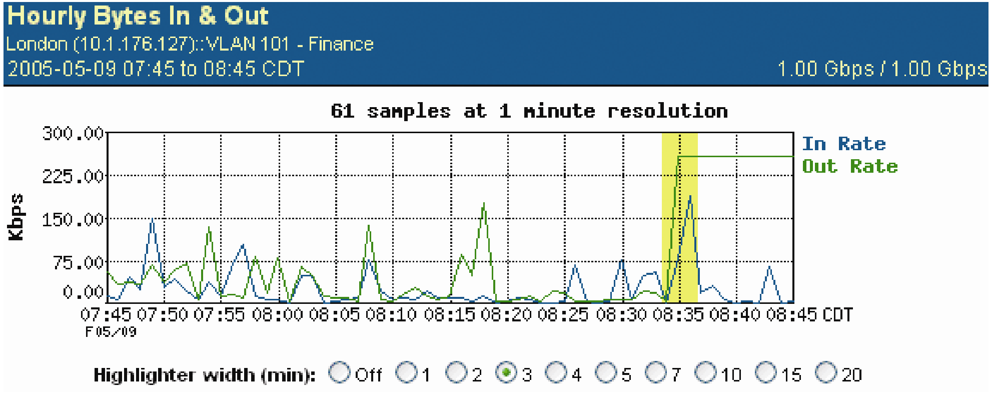

Hourly Bytes In & Out
London (10.1.176.127)::VLAN 101 - Finance
2005-05-09 07:45 to 08:45 CDT 1.00 Gbps / 1.00 Gbps

61 samples at 1 minute resolution

Kbps
300.00
225.00
150.00
75.00
0.00
07:45 07:50 07:55 08:00 08:05 08:10 08:15 08:20 08:25 08:30 08:35 08:40 08:45 CDT
F05/09
In Rate
Out Rate

Highlighter width (min): ○ Off ○ 1 ○ 2 ● 3 ○ 4 ○ 5 ○ 7 ○ 10 ○ 15 ○ 20

The engineer sees the spike in rate across the interface occurring 10 minutes ago and highlights it to investigate further. The detail report is shown in Figure 7.

Figure 7

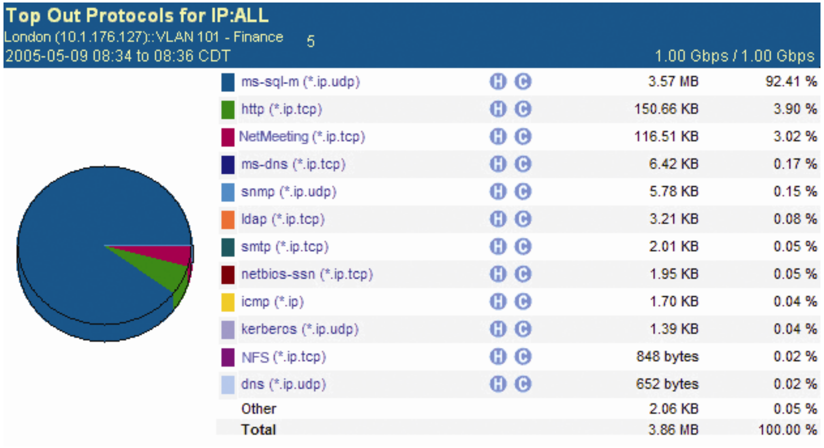

Top Out Protocols for IP:ALL
London (10.1.176.127)::VLAN 101 - Finance 5
2005-05-09 08:34 to 08:36 CDT 1.00 Gbps / 1.00 Gbps

🟦 ms-sql-m (*.ip.udp)	Ⓗ Ⓒ	3.57 MB	92.41 %
🟩 http (*.ip.tcp)	Ⓗ Ⓒ	150.66 KB	3.90 %
🟥 NetMeeting (*.ip.tcp)	Ⓗ Ⓒ	116.51 KB	3.02 %
🟦 ms-dns (*.ip.tcp)	Ⓗ Ⓒ	6.42 KB	0.17 %
🟦 snmp (*.ip.udp)	Ⓗ Ⓒ	5.78 KB	0.15 %
🟧 ldap (*.ip.tcp)	Ⓗ Ⓒ	3.21 KB	0.08 %
🟩 smtp (*.ip.tcp)	Ⓗ Ⓒ	2.01 KB	0.05 %
🟥 netbios-ssn (*.ip.tcp)	Ⓗ Ⓒ	1.95 KB	0.05 %
🟨 icmp (*.ip)	Ⓗ Ⓒ	1.70 KB	0.04 %
🟪 kerberos (*.ip.udp)	Ⓗ Ⓒ	1.39 KB	0.04 %
🟪 NFS (*.ip.tcp)	Ⓗ Ⓒ	848 bytes	0.02 %
🟦 dns (*.ip.udp)	Ⓗ Ⓒ	652 bytes	0.02 %
Other		2.06 KB	0.05 %
Total		3.86 MB	100.00 %

The most used protocol during the spike is ms-sql and the engineer has just realized that the SQL Slammer virus is active on the network. The network engineer now wants a comprehensive list of hosts that are infected with the virus. Using the Flow Forensics Wizard, he configures an Application Client report [Figure 8] to see a comprehensive list of hosts passing that protocol above a specified flow rate for the past 4 hours. With Flow Forensics, the engineer can see 100% of the flow traffic on the network with no byte thresholds or top-n reporting requirements.

Figure 8

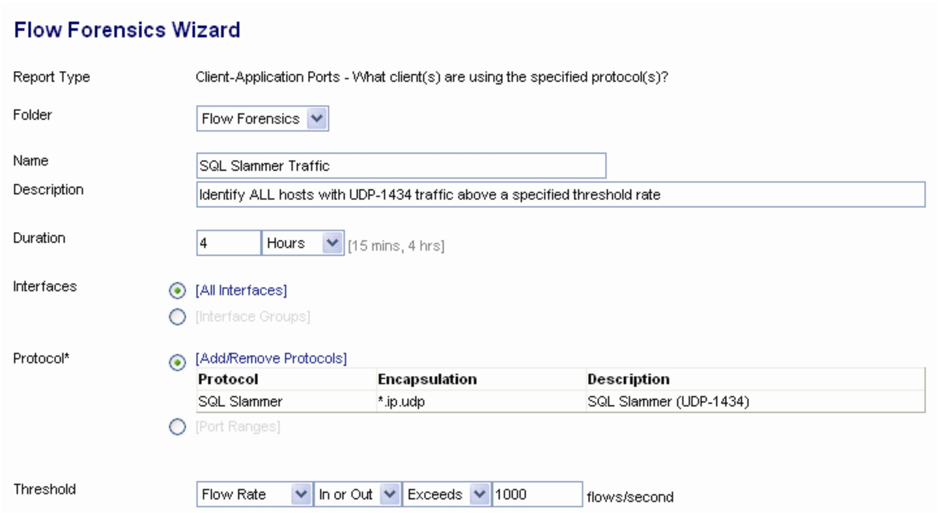

Flow Forensics Wizard

Report Type: Client-Application Ports - What client(s) are using the specified protocol(s)?

Folder: Flow Forensics

Name: SQL Slammer Traffic

Description: Identify ALL hosts with UDP-1434 traffic above a specified threshold rate

Duration: 4 Hours [15 mins, 4 hrs]

Interfaces: ● [All Interfaces] ○ [Interface Groups]

Protocol*: ● [Add/Remove Protocols]

Protocol	Encapsulation	Description
SQL Slammer	*.ip.udp	SQL Slammer (UDP-1434)

○ [Port Ranges]

Threshold: Flow Rate | In or Out | Exceeds | 1000 flows/second

In seconds, the engineer has a report that identifies all the hosts infected with the virus, shown in Figure 9.

Figure 9

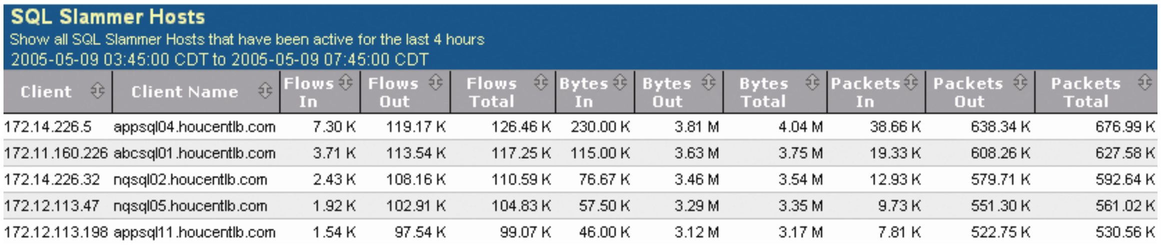

SQL Slammer Hosts

Show all SQL Slammer Hosts that have been active for the last 4 hours

2005-05-09 03:45:00 CDT to 2005-05-09 07:45:00 CDT

Client	Client Name	Flows In	Flows Out	Flows Total	Bytes In	Bytes Out	Bytes Total	Packets In	Packets Out	Packets Total
172.14.226.5	appsql04.houcentlb.com	7.30 K	119.17 K	126.46 K	230.00 K	3.81 M	4.04 M	38.66 K	638.34 K	676.99 K
172.11.160.226	abcsql01.houcentlb.com	3.71 K	113.54 K	117.25 K	115.00 K	3.63 M	3.75 M	19.33 K	608.26 K	627.58 K
172.14.226.32	nqsql02.houcentlb.com	2.43 K	108.16 K	110.59 K	76.67 K	3.46 M	3.54 M	12.93 K	579.71 K	592.64 K
172.12.113.47	nqsql05.houcentlb.com	1.92 K	102.91 K	104.83 K	57.50 K	3.29 M	3.35 M	9.73 K	551.30 K	561.02 K
172.12.113.198	appsql11.houcentlb.com	1.54 K	97.54 K	99.07 K	46.00 K	3.12 M	3.17 M	7.81 K	522.75 K	530.56 K

Solution: In this example, ReporterAnalyzer allowed the engineer to access a report showing traffic at one minute granularity for the last hour to identify what was on the network in real time. The protocol causing inaccessibility for the application resources was a known virus and all hosts above a specified flow rate of the protocol were identified. The engineer was able to take the infected hosts offline and apply the necessary virus updates and patches to prevent the SQL Slammer virus from propagating further across the network.

Example 3: Optimizing the Infrastructure for Application Performance: What is causing the poor VoIP performance?

Problem: Having performed all the necessary network assessments, a company rolls out VoIP to a remote site in Houston with the confidence that their infrastructure can handle the load. Unfortunately, end users begin to complain about poor, and sometimes unavailable, VoIP service.

Analysis: The network engineer receiving the calls needs access to detailed information about traffic occurring on the Houston link to diagnose the problem. The engineer views the protocol distribution on the network link, and sees that voice traffic is well within an acceptable range. See Figure 10.

Figure 10

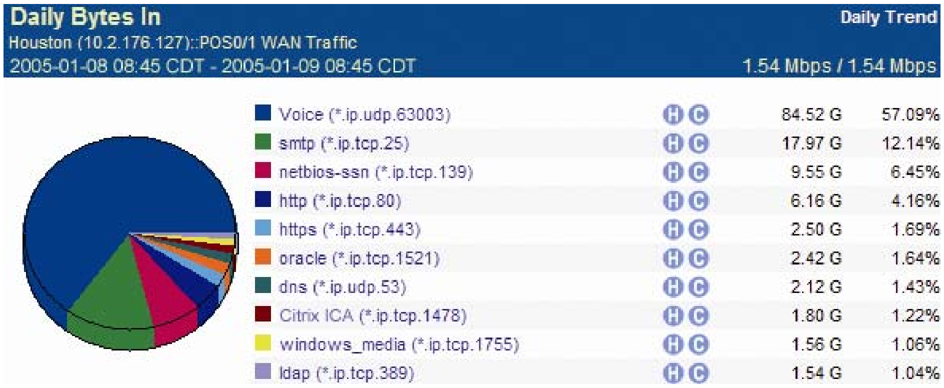

Daily Bytes In

Houston (10.2.176.127)::POS0/1 WAN Traffic

2005-01-08 08:45 CDT - 2005-01-09 08:45 CDT

■	Voice (*.ip.udp.63003)	H C	84.52 G	57.09%
■	smtp (*.ip.tcp.25)	H C	17.97 G	12.14%
■	netbios-ssn (*.ip.tcp.139)	H C	9.55 G	6.45%
■	http (*.ip.tcp.80)	H C	6.16 G	4.16%
■	https (*.ip.tcp.443)	H C	2.50 G	1.69%
■	oracle (*.ip.tcp.1521)	H C	2.42 G	1.64%
■	dns (*.ip.udp.53)	H C	2.12 G	1.43%
■	Citrix ICA (*.ip.tcp.1478)	H C	1.80 G	1.22%
■	windows_media (*.ip.tcp.1755)	H C	1.56 G	1.06%
■	ldap (*.ip.tcp.389)	H C	1.54 G	1.04%

As part of the VoIP rollout plan, voice traffic was assigned a "Platinum" quality of service level. This ToS value would ensure voice to have the highest priority within the network.

The engineer begins to suspect that the ToS configuration for voice traffic was improperly configured, so he views the protocol summary for the "Platinum" ToS level. See **Figure 11**.

Figure 11

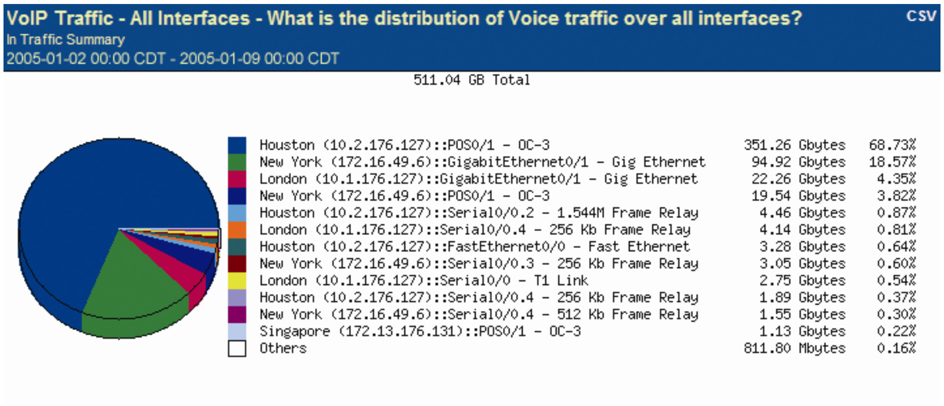

VoIP Traffic - All Interfaces - What is the distribution of Voice traffic over all interfaces?

In Traffic Summary

2005-01-02 00:00 CDT - 2005-01-09 00:00 CDT

511.04 GB Total

🟦	Houston (10.2.176.127)::POS0/1 - OC-3	351.26 Gbytes	68.73%
🟩	New York (172.16.49.6)::GigabitEthernet0/1 - Gig Ethernet	94.92 Gbytes	18.57%
🟥	London (10.1.176.127)::GigabitEthernet0/1 - Gig Ethernet	22.26 Gbytes	4.35%
🟦	New York (172.16.49.6)::POS0/1 - OC-3	19.54 Gbytes	3.82%
🟦	Houston (10.2.176.127)::Serial0/0.2 - 1.544M Frame Relay	4.46 Gbytes	0.87%
🟧	London (10.1.176.127)::Serial0/0.4 - 256 Kb Frame Relay	4.14 Gbytes	0.81%
🟩	Houston (10.2.176.127)::FastEthernet0/0 - Fast Ethernet	3.28 Gbytes	0.64%
🟫	New York (172.16.49.6)::Serial0/0.3 - 256 Kb Frame Relay	3.05 Gbytes	0.60%
🟨	London (10.1.176.127)::Serial0/0 - T1 Link	2.75 Gbytes	0.54%
🟪	Houston (10.2.176.127)::Serial0/0.4 - 256 Kb Frame Relay	1.89 Gbytes	0.37%
🟪	New York (172.16.49.6)::Serial0/0.4 - 512 Kb Frame Relay	1.55 Gbytes	0.30%
🟦	Singapore (172.13.176.131)::POS0/1 - OC-3	1.13 Gbytes	0.22%
□	Others	811.80 Mbytes	0.16%

In reality, voice traffic was properly configured for "Platinum" level of service, but the CRM application for the company was responsible for the poor VoIP performance. Lowering the level of service for the CRM application to "Gold" dramatically improved performance to the degree they had anticipated.

Solution: In this example, ReporterAnalyzer helped the engineer discover other improperly configured application ToS levels that directly impacted VoIP usage. The configuration error was corrected and performance of the voice application on the link returned to normal.

The NetQoS NetFlow/IPFIX Solution Summary

NetFlow/IPFIX and ReporterAnalyzer are powerful toolsets that can be used together to help network engineers confidently manage complex networks. Information from these tools can be used to decide necessary changes in infrastructure for capacity analysis and planning, to troubleshoot network problems in real time, to analyze network traffic for billing purposes, and to reposition IT staff from reactive to proactive network management.

"NetQoS captures the NetFlow export data from Cisco routers and switches that already exist in many enterprises today, providing a viable, scalable solution to network congestion. We selected NetQoS as our capacity planning application, largely because it took advantage of the Cisco IOS NetFlow capability present on all Cisco routers to analyze traffic and detect problems."

- *Cisco*The NetQoS ReporterAnalyzer solution enables enterprise customers to leverage their existing Cisco routers and switches by using the NetFlow/IPFIX data already available from them.

Conclusion

This whitepaper has described two possible methods of collecting network traffic analysis data – using Cisco NetFlow and using RMON2 probes. Each enterprise should evaluate the options available and make an informed decision based on the business objectives and resources available.

About the NetQoS Performance Center

The NetQoS Performance Center is a single web-based portal that delivers global visibility into the entire network infrastructure for the insight to resolve performance issues, troubleshoot infrastructure problems, perform capacity planning, plan for and validate the effects of change, and track Service Level Agreements (SLAs). The NetQoS Performance Center provides an integrated view of critical performance data delivered by NetQoS' best in class products—SuperAgent® for end-to-end performance data, ReporterAnalyzer™ for traffic analysis, and NetVoyant® for device performance (SNMP Polling) data. With this visibility, you will make more informed decisions based on precise network infrastructure usage data, improve staff efficiency, and resolve performance issues rapidly. You can also mitigate risks and validate the impact of planned changes such as VoIP deployments, MPLS migrations, WAN optimization, QoS policy implementation, and application rollouts. The NetQoS Performance Center is fast and easy to deploy, and scales to the largest networks.

About NetQoS

NetQoS is the fastest growing network performance management products and services provider. NetQoS has enabled hundreds of the world's largest organizations to take a Performance First approach to network management—the new vanguard in ensuring optimal application delivery across the WAN. By focusing on the performance of key applications running over the network and identifying where there is opportunity for improvement, IT organizations can make more informed infrastructure investments and resolve problems that impact the business. Today, NetQoS is the only vendor that can provide global visibility for the world's largest enterprises into all key metrics necessary to take a Performance First management approach. More information is available at **www.netqos.com**.

NetQoS, Inc.
e. info@netqos.com
p. 512.407.9443
t. 877.835.9575
f. 512.407.8629
www.netqos.com