InMon Corp.
1404 Irving St., San Francisco, CA 94122
www.inmon.com

Traffic Monitoring in a Switched Environment

1. SUMMARY

This document provides a brief overview of some of the issues involved in monitoring traffic in a switched environment. The cost, scalability and functionality of currently available technologies are compared. Finally, a table is provided that matches requirements to appropriate technologies.

2. BACKGROUND

Traffic monitoring is a vital element of network and system management. Very little happens in an enterprise without producing some network traffic. Monitoring this traffic gives important information about the operation of enterprise applications. This information is essential for activities such as cost allocation, capacity planning, quality of service analysis, fault detection and isolation and security management.

Traffic monitoring used to be a relatively straightforward task. In the past large numbers of machines were connected to a shared network. A shared network permits a single instrument connected to the network to monitor all the traffic since packets sent in one part of the network are received in all other parts of the network (see Figure 1).

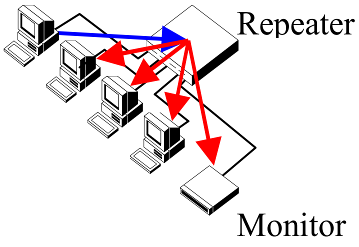

Figure 1: ***Shared Network***

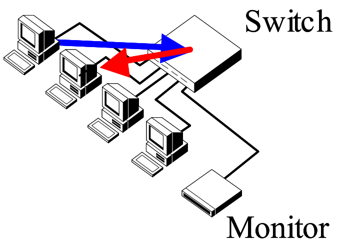

Figure 2: ***Segmented Network***

Requirements for increased bandwidth, changes in traffic patterns, and the quickly falling price of packet switching and routing devices has caused a rapid movement away from shared networks to networks which are highly segmented (see Figure 2).

Traffic is no longer visible from a single point. A switch directs packets to specific ports based on the packet's destination. Every port on the switch needs to be monitored in order to obtain a complete picture of the network traffic. The use of point-to-point links makes it difficult to attach instruments and the large number of instruments that would be required to monitor all the switch ports ensures that such an approach would not be cost effective. In addition the switches and routers themselves have complex internal architectures and the flow of packets within, and through, them is becoming an important factor in network performance.

The only realistic way to monitor traffic on switched networks is to monitor traffic within the switches themselves. In addition to the technical difficulties of the task there are also severe price constraints. The market for switches is maturing and there is very little room to add cost or impact the performance of these devices, especially since monitoring is secondary to the primary switching function of the device.

3. CURRENT SOLUTIONS AND LIMITATIONS

There are a number of approaches to monitoring network traffic, each of which has different strengths and weaknesses.

Currently there are three main choices for traffic monitoring:

RMON RMON[1] (Remote MONitor) is an Internet Engineering Task Force (IETF) standard specifying a remote, promiscuous, traffic-monitoring device. An RMON device monitors and decodes every packet on the network to which it is attached, creates tables of measurements that can be later downloaded by a network management application.

NetFlow® Cisco routers and switches, as part of their NetFlow® monitoring system, send information about completed traffic flows, to a central collector. The device decodes every IP packet, maintains tables of active flows, and forwards flow records periodically or when they complete to a network management application.

sFlow™ sFlow combines accurate packet counters with a statistical sampling of the state of the routing and bridging tables used by the switch to forward randomly selected packets. The sampled information is immediately sent to a central collector for analysis.

The following sections compare RMON, NetFlow® and sFlow in a number of key areas.

A. Agent Resources

When comparing technologies for use in an embedded monitoring application one must consider the three main resources that any embedded traffic monitoring solution will consume: CPU, memory and bandwidth. These resources are expensive and an embedded measurement solution must minimize resource consumption in order to be cost effective. The following charts compare RMON, NetFlow® and sFlow in each of the three resources.

a. CPU

The computational requirements of traffic monitoring significantly impact agent costs and scalability. Computationally intensive monitoring techniques require a high performance network management CPU in a switch or router, adding to its cost. In addition a single processor may not be able to cope with the demands of monitoring the large numbers of ports contained in many switches and so additional CPU's may be required for monitoring. Figure 3 compares the CPU load needed to monitor traffic using each of the different technologies.

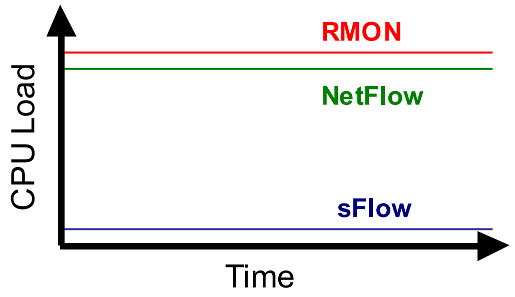

RMON : high
Decoding every packet and inserting in table, while handling polling requests and encoding data in ASN.1 for SNMP upload)

NetFlow : high
Decoding every packet and inserting in table.

sFlow : very low
Can be implemented in hardware.

Figure 3 ***Comparison of Agent CPU Load***

Both RMON and NetFlow® agents attempt to build traffic matrices by decoding every packet. Given that a 100Mb/s link can theoretically exceed 200,000 packets/second, this places a worst-case load that would saturate all but the fastest CPUs today. Even a load of only 10,000 packets/second means a packet must be processed every 100 microseconds. (If the CPU cannot keep up, packets are missed. Thus the results can be systematically biased.)

[1] S. Waldbusser, "Remote Network Monitoring Management Information Base, Version 2", RFC 2021.

Constructing a traffic matrix in agent memory creates a variable, unpredictable load on the CPU, depending on the nature of the traffic patterns. Both RMON and NetFlow suffer from this variable overhead.

In the case of RMON, the CPU also has to enter into a dialog with a network management station into order to upload the results from the last period. Marshalling data into ASN.1 format for SNMP upload incurs a further significant overhead.

In contrast, the sFlow agent's sampling function can easily be implemented in hardware. Thus if the network was bursting at 200,000 packets/second and the sampling rate was 1/1000 the agent CPU would only be required to handle 200 packets/second. Since it is not required to do anything but forward the packet header to the server, this load is very light indeed.

b. Memory

The amount of memory required to construct traffic measurements affects the cost of the agent. In addition, variable memory requirements cause special problems in deciding how much memory to incorporate into a device. Too little memory and the traffic monitor will run out of memory and loose data, too much memory adds unnecessary cost to the agent. Figure 4 compares the different monitoring technologies in terms of the amount and variability of their agent memory requirements.

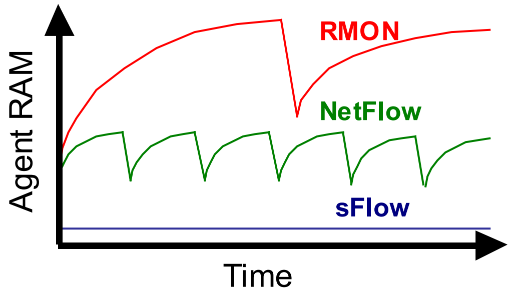

RMON : very high (8-32MB)
Table grows continuously until retrieved.

NetFlow : high (4-8MB)
Tables flushed automatically.

sFlow : very low (1KB)
Space for one sample packet.

Figure 4 ***Comparison of Agent Memory Usage***

Both RMON and NetFlow® agents build traffic matrices in agent RAM. The size of the tables is heavily dependent on the traffic patterns. In the worst case, every packet can represent a new flow for which addresses and counters must be stored separately. NetFlow® has the benefit that it can be configured to flush individual flows automatically every 15 minutes or so. A NetFlow® agent is also able to flush blocks of flows to prevent it running out of memory. However, an RMON agent must keep the entire traffic matrix in memory and start a new one for the next period while it waits for the server to come and retrieve the results. Additional RMON measurements, such as for top talkers, all require still more agent memory.

The sFlow agent needs only enough RAM to store one packet. When a sample is taken, it is forwarded immediately to the server.

c. Bandwidth

Transferring measurements from monitoring agents to a central collector for reporting and analysis can consume significant bandwidth. Bandwidth used to monitor the network is bandwidth taken away from network applications.

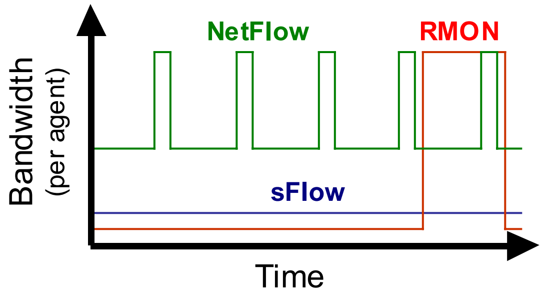

RMON : bursty
Low for stats polling, but full-rate burst for SNMP data retrieval.

NetFlow : high bursty
Full rate burst during flush.

sFlow : low smooth
Typically about 1 packet per second, can be controlled by adjusting sampling rate.

Figure 5 ***Comparison of Bandwidth***

The burstyness of transfers is also important - high peak transfer rates can cause significant congestion in the network. Figure 5 compares the different monitoring technologies in terms of network bandwidth.

The network load from a NetFlow® agent appears as a stream of large UDP packets. Usually these contain flows that are known to be complete or are due for their periodic flush, but as the flow table becomes full, large numbers of flows may be flushed at once in order to free space. An agent on a moderately loaded link can generate an average of about 30 packets/second.

An RMON agent collects the entire traffic matrix in memory, with no incremental flush. At the end of the collection period (e.g. 1 hour) there is a burst of network activity as the server retrieves it. This burst creates a scheduling problem for the server: at the end of the hour it wants to collect the last hour's traffic data from every agent in the network, preferably as soon as possible because the agent RAM may be filling up with the next hour's data. This limits the number of agents that can be managed with a single server.

Any additional RMON measurements incur a further increase in the network traffic as the server must set up each measurement and retrieve its results separately.

The sFlow agent's sampling process results in a steady stream of packets from all the agents to the server. Simply adjusting the sampling rate each agent uses can control the traffic level.

B. Server Resources

The traffic-monitoring agent is only a small part of an overall traffic management solution. A central traffic server is needed to configure the agents, download and archive data, and make the data available through a user interface. A key factor in determining the cost and scalability of the traffic monitoring system is to determine server resources needed to handle a single agent.

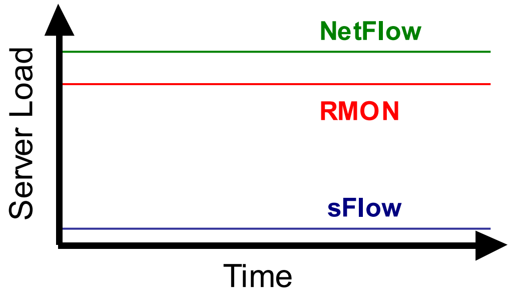

RMON : high
Stats polling and scheduled data collection using SNMP requires almost as much traffic out as traffic in. Consolidating site-wide traffic-matrices is a heavy task because un-sampled traffic matrices can be large

NetFlow : high
No traffic out, but more traffic in because traffic matrices are more detailed than RMON. This also makes the consolidation task even heavier.

sFlow : moderate
Low traffic out, sampling reduces size of traffic matrices..

Figure 6 ***Comparison of Server Load***

The resource consumption at the server determines how many agents a single server can manage and significantly affects the cost and scalability of the overall traffic monitoring system. Figure 6 compares each technology in terms of the server resources needed to handle an agent.

RMON and NetFlow® servers both suffer from the large size of their agent traffic matrices. Apart from the effort of receiving them over the network, there is still the task of consolidating them into site traffic matrices (taking care not to count the same flow more than once if it was seen by more than one agent).

For RMON, polling for segment counters and collecting top-talkers data represents a significant additional load.

The size of a traffic matrix constructed from sampled packets is significantly smaller. Typically it consists of the busiest flows, plus a random selection of the smaller flows (many flows in the network consist of only 4 or 5 packets). The task of consolidating these traffic matrices is therefore much lighter. The server also has the benefit that the segment counters are included with the samples, and top-talkers can be calculated from the same samples too.

C. Features

A comparison of technologies purely in terms of scalability and resource consumption would be meaningless. It is essential to also compare the functionality provided by each system. The following three areas represent key functional areas that a traffic monitoring system should address.

a. Real-Time Segment Statistics

Tracking and trending the segment statistics for every link in the network is fundamental to good network management.

RMON	Supported - but must be polled from server using SNMP. This limits the number of agents that can be monitored.
NetFlow®	Outside scope of specification. Can use SNMP to obtain segment counters from most devices.
sFlow	Supported - available from every agent all the time with no requirement for polling (interface counters are included with samples from agent).

For an RMON server to collect minute-by-minute segment statistics measurements it must use SNMP to poll all the agents in the network at least every minute. Synchronizing all these measurements to begin and end on the same minute boundary is clearly impossible. The polling must be spread out over time. The overhead of SNMP limits the number of counters that can be requested in one packet. The extra traffic generated can be significant. Most RMON solutions react by pushing the alarm thresholding down onto the agent, and thus lose the ability to trend and correlate over time and between segments.

The NetFlow® specification does not address the issue of segment statistics. Typically SNMP can be used to obtain segment statistics. The issues in using SNMP to obtain segment statistics are described in the paragraph on RMON.

With the sFlow agent, the hardware segment counters from each agent are piggybacked onto the sample stream with a space-efficient encoding. Thus the server can see up-to-date counters from every segment without having to request them separately. Full segment statistics can be computed for every segment for every minute - synchronized to the same global minute boundary.

b. Real-Time Top Talkers

It is not enough to know that a segment is congested. To do anything about it, it is necessary to know why. Top talkers measurements provide the answer.

RMON Supported as point study - typically 1 measurement on one agent at a time. Incurs traffic overhead as results must be retrieved using SNMP. This limits the number of agents that can be monitored.

NetFlow® Not supported.

sFlow Top sources, destinations and pairs for every protocol, every minute on every agent and port - continuously. No traffic overhead as all results are derived from incoming samples.

With RMON, collecting top-talkers information from an agent requires the measurement to be set-up and the results to be retrieved. Doing this every minute for every segment on the network represents a huge task for the server. Furthermore, computing the top talkers in an RMON agent is very expensive in terms of both CPU and RAM. In an RMON system, top-talkers can only realistically be used as a point-study (perhaps as a follow-up measurement when a segment counter trips a threshold).

NetFlow® does not support real-time top talkers measurements. The delay in transferring flow records from the agent to the server makes it impossible for the server to provide accurate real-time data.

With sFlow the top talkers are computed on the server from the same stream of samples that is used to build the traffic matrices. Computing top talkers from sampled data is much easier than when looking at every packet. The server can offer top sources, top destinations and top talking pairs for every protocol for every minute for every agent and port. Extra measurements can be added, such as top sources of IP multicasts, without any increase in the network or agent load. The server can even correlate across thousands of segments and trend the top flows within a site or subnet.

c. Traffic Matrices

Site-wide traffic matrices are important for capacity planning, billing, topology optimization and reporting. They represent a snapshot of the “offered” load on the network that is independent of how the underlying topology is actually carrying it. Thus the data remains relevant and useful even if nodes are moved and the topology is changed.

RMON Extensive Suite (limited by agent).

NetFlow® IP, ICMP, TCP, UDP (limited by agent)

sFlow Extensive Suite (not limited by agent because decodes are performed on server. Hence new traffic matrices can be added any time without having to upgrade the agents).

The RMON standard offers an extensive suite of traffic matrices, although it is doubtful that an RMON agent would have enough RAM to build all of them at once (even if it did, there is still the problem of transferring the data to the server for consolidation). Another drawback is that the agents have to be upgraded if a new protocol decode is to be supported.

NetFlow® supports only IP, ICMP, TCP and UDP.

With sFlow, decodes are performed on the server and there is no extra network or agent load if multiple traffic matrices are calculated.

4. CONCLUSIONS

Table 1 summarizes the different monitoring technologies in terms of their scalability and the applications they are best suited to.

	RMON	NetFlow®	sFlow
Scalability (Maximum number of ports per server).	~ 100	~ 10	~ 50,000
Recommended Use	Remote protocol analyzer, deployed for troubleshooting.	Billing and security monitoring of moderate speed WAN links on routers.	Monitoring of all switch ports. High-speed backbone link monitoring. Billing, congestion management and security.

Table 1 ***Comparison of Traffic Monitoring Technologies***

When collecting minute-by-minute segment statistics and top talkers, plus site-wide traffic matrix consolidation, it is difficult to see how any RMON server could support more than 100 agents simultaneously. With the high unit cost of RMON agent technology, even this rather low number adds up to a prohibitively expensive solution.

What the RMON standard defines is essentially a remote protocol analyzer. It's ability to filter, capture and decode packets make it applicable to those troubleshooting occasions where a problem can only be understood by seeing the timestamped protocol sequences on the wire.

The NetFlow® technology can be useful on WAN links where it may be essential to see every flow, perhaps for security monitoring or specific flow-by-flow billing. The sheer volume of data generated means that a server cannot be expected to manage more than 10 links.

The statistical sampling technique employed by sFlow scales well to large numbers of agents, tens of thousands of switch ports can be managed by one server. Its zero agent cost and significant technical advantages make it ideal for continuous site-wide (and enterprise-wide) traffic monitoring and reporting.

5. MORE INFORMATION

For more information contact InMon Corp. InMon has extensive experience in implementing sampling-based traffic management. InMon's Traffic Server provides means of continuously monitoring and reporting on traffic in large switched environments.

InMon Corp.
1404 Irving Street
San Francisco, CA 94122

info@inmon.com http://www.inmon.com/

(415) 661-6343