INMON CORP.
580 CALIFORNIA ST., 5th FLOOR,
SAN FRANCISCO, CA 94104
WWW.INMON.COM

sFlow® Accuracy and Billing

1 Introduction

There is a growing interest in an accurate and effective way to capture network and network service usage data in order to generate itemized bills, either for generating revenue or for apportioning network service costs. Generating itemized bills is an important value-added service that has two major advantages:

1. A customer has confidence that the bill accurately reflects their usage, and provides them with information that can be used to control costs.
2. The service provider can charge for differentiated services, creating new revenue opportunities.

Usage data that supports itemized billing must give accurate information on volume of traffic transmitted and received. The method for collecting usage data must be inexpensive and should not impact the performance of the network.

This application note describes how the sFlow technology can meet these goals.

2 sFlow overview

sFlow was designed to provide continuous network-wide monitoring of L2-L7 packet flows in high-speed switched environments. A key element of sFlow is the statistical sampling of packets forwarded by a switch or router. Once a packet has been sampled, the packet header (often the first 128B) is forwarded to a central data collector (e.g. InMon Traffic Server) that performs analysis on samples collected from a complete network of switches and routers. This results in detailed information on application level traffic flows. This data is used to generate itemized usage breakdowns on many levels, for example: by VLAN, priority, type of service, local and non-local traffic.

3 Sampling theory

The sFlow statistical sampling is a count-based (or packet-based) sampling technique. On average, one packet in N is sampled and forwarded for analysis. An element of randomness is introduced into the sampling process to prevent synchronization with any periodic patterns in the traffic.

Sampling does not provide a 100% accurate result, but it does provide a result in which the error can be accurately characterized.

http://www.sflow.org/packetSamplingBasics/index.htm describes the statistical theory behind sFlow.

This theory is best illustrated by example.

3.1 Packet sampling example

Suppose 1,000,000 packets are forwarded by a switch and a random sample of 0.25% is taken (2,500 packets). If 1,000 of the samples represent a particular class of traffic (voice traffic say), then how many packets crossing the switch were actually voice packets?

The most likely fraction of voice packets is the same as the fraction of samples that were voice packets i.e. 40% (1,000 divided by 2,500). This means that the best estimate for number of voice packets is 400,000.

It is unlikely that there were exactly 400,000 voice packets. Instead it is possible to specify, mathematically, a small range of values that are very likely, say 95% likely. In other words, a range, or 95% confidence interval, can be calculated mathematically, and one can be 95% confidant that the actual number of voice packets falls in this range. In this case the 95% confidence interval is between 381,000 and 419,000.

This range can also be expressed as a percentage of the most likely value, i.e. 400,000 ±4.8% were voice packets or the largest likely error is 4.8%.

An HP Labs Technical Report http://www.hpl.hp.com/techreports/92/HPL-92-35.html describes how to compute confidence intervals from sampled data. The %error, to a first approximation, is a function of only the number of samples used to make the measurement:

$\%\text{error} \le 196 * \sqrt{\frac{1}{c}}$ where c is the number of packets that belong to the class.

For this example c was 1,000, which gives an error of approximately 6%.

4 Tuning sFlow for billing

In a typical billing/accounting application, the objective is to determine from all the packets crossing the network during the billing period (often 1 month), how many packets came from a particular source. For example, a common ISP billing strategy is a flat monthly charge that includes an allowance (say 4 GB) with additional charges for each additional GB.

In an environment in which the sampling rate is set to 1 in 1000, approximately 33,000 samples will be taken from traffic sent by a source which transmits 5 GB during a month with average packet size of 150B. Using the equation above, the %error would be within 1.07%.

The accuracy of the sFlow results can be tuned to achieve the desired level of accuracy appropriate to the billing strategy. Since the accuracy is related to the number of samples received, the accuracy can be tuned by increasing the number of samples, this can be done in two ways:

1. Increasing the sampling rate
2. Monitoring over a longer period of time

The graph in Figure 1 is useful for determining the number of samples that should be collected to obtain a given % error. This can then be used to determine the sampling rate.

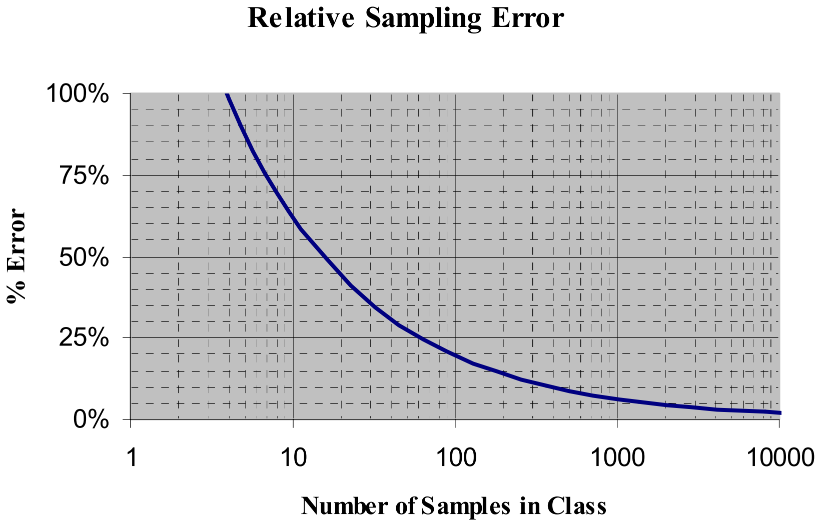

Figure 1 Relative sampling error

In addition the confidence intervals for measurements can be used to fairly address the issue of errors for the customer. Billing by lower bound of the confidence interval ensures that no customer is overcharged.

Because the accuracy of the sFlow measurement is dependent only on the number of samples, if samples are lost between the sFlow agent, performing the sampling in the device, and the central data collector, the increase in %error in the result is negligible and is easily calculated.

5 Other monitoring methodologies for billing

There are two other commonly used methodologies for collecting traffic data for billing: polling counters and examining every packet to derive flow information.

5.1 Polling counters

Basic frame and byte counters maintained by each network device interface can be polled. This gives accurate aggregate counts of frames and bytes passing through the interface.

If a single host is attached to a switch port, the total volume of traffic to that host will be measured. If multiple services are hosted on that machine (e.g. multiple customers' web sites), then it is not possible to distinguish between traffic destined for the different customers. Similarly, it is not possible to distinguish between Internet web traffic and local backup traffic and therefore bill differently for local and non-local traffic.

In practice errors may occur with this methodology, for example counters may wrap or line cards may be reset. Because these situations are unpredictable, it is not possible to quantify or characterize the errors.

sFlow also uses accurate packet counters, forwarding them with samples to the central data collector. However, because sFlow is implemented within the switch, it is able to ensure that counter wraps and resets are detected and do not result in measurement errors.

5.2 Aggregating traffic flow data on a switch

NetFlow operates by accumulating traffic flow totals into an onboard flow cache. For the most detailed data, totals are accumulated for unique tuples (flow) of source IP address, source UDP/TCP port, destination IP address, destination UDP/TCP port. This method requires a variable, but significant amount of memory, especially under high load conditions, for example during a denial of service attack when every packet is a separate short-lived flow there may be 30,000 flows per second and the switch must export data rapidly to avoid flow cache overflow. In such a situation flow data will be lost.

Because aggregation is performed on the device before being sent to a central data collector, a single measurement can represent a significant fraction of the overall traffic, if a packet containing this data is lost, the accuracy of the overall measurement will be impacted.

The impact of these situations is impossible to quantify, and therefore the final accuracy of the measurement cannot be characterized.

6 Conclusion

Packet-based sampling is an inexpensive monitoring technology that easily keeps up with today's high-speed switched networks. As implemented by sFlow, there are a number of key advantages:

- Detailed characterization of packet flows that supports itemized bills on many levels e.g. VLAN, protocol, type of service, source and destination.
- Error bounds are easily calculated and monitoring can be tuned to achieve the desired accuracy.
- Robust in the face of loss of packets containing measurement data.
- Does not impact the performance of the network devices.