

Benefits of Flow Analysis Using sFlow[®]: Network Visibility, Security and Integrity

**3650 Brookside Pkwy
Brookside Concourse 100
Suite 400
Alpharetta, Georgia 30022
P: 770.225.6500
F: 770.225.6501**

**SALES@LANCOPE.COM
WWW.LANCOPE.COM**

Lancope.[®]
>security through network intelligence™

SUMMARY	3
Introduction to sFlow	3
Introduction to Flow Processing	4
Flow Processing using sFlow	4
Using Sampled Data for Security	4
Protocol Violations.....	5
“Dark IP Space” Access Detection	5
User Policy Violations.....	5
Peer-to-Peer (P2P) Application Detection.....	5
New Host Identification.....	5
Operating System Identification	5
Denial of Service (DoS).....	5
Scanning “Attacks”	6
Summary of Sampled Data for Security	8
Traffic Analysis	8
sFlow Sampling Strategies	8
Case 1: All Edge Switches Are sFlow Capable.....	9
Case 2: Some Edge Switches Are sFlow Capable	9
Case 3: Only Core Switches Are sFlow Capable.....	9
Recommended Configuration.....	10
sFlow Bandwidth	10
Monitoring High Speed Network Links	10
StealthWatch System	10
StealthWatch for sFlow: Competitive Differentiators	11
NBA Solutions using sFlow	11
Conclusion	11
About Foundry® Networks	12
About Lancopé®	12
About the Author	12

Legal Notices and Disclaimers: The information contained in this document is proprietary and confidential to Lancopé. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose without the express written permission of Lancopé. For information on site licenses and multiple copy discounts, contact Lancopé. This document is subject to change without notice. While Lancopé has endeavored to provide a high level of accuracy, no complete assurances of accuracy can be provided. If you find any problems with this document, please report them to Lancopé in writing.

Lancopé is a registered trademark and StealthWatch™ is a trademark of Lancopé, Inc. sFlow is a registered trademark of InMon, Corp. Foundry Networks and the ‘Iron’ family of marks are trademarks of Foundry Networks, Inc. All other registered or unregistered trademarks are the sole property of their respective owners.

© 2006 Lancopé, Inc. All rights reserved.

Summary

This technical white paper provides the reader with an introduction sFlow® and the network visibility, integrity and security benefits it provides for wide area and very high bandwidth network monitoring. The paper discusses how detailed statistical analysis of sampled packets can provide a comprehensive network view, encompassing both traffic analysis and security, when combined with an innovative comprehensive Network Behavior Analysis (NBA) and Response technology, such as StealthWatch™ by Lancope®. Included are several examples that highlight how various network attacks are detected as compared to traditional packet capture and NetFlow-based solutions.

Introduction to sFlow

sFlow, a technology designed for network monitoring based on packet sampling, derives from early work performed by HP at the University of Geneva and CERN in 1991. This was followed by the eventual definition of sFlow as RFC-3176 in 2001.ⁱ Since then, development on sFlow has continued with additional extensions to the RFC being developed and supported by major manufacturers such as Foundry Networks, Inc.

sFlow operates by randomly sampling one out of every “n” packets at an observation point on the network – typically in a switch or router. The packet header along with forwarding tables and interface counters are provided to the sFlow agent embedded in the device, which then forwards the information to sFlow collectors for processing.

The basic theory of packet sampling for network traffic analysis can be modeled by Binomial Distribution. The theory is based on obtaining a significantly large number of samples over a period of time to minimize the sampling error for a “class” of traffic (all packets, HTTP packets, packets to IP address 10.1.1.1, etc.). The error rate is approximated by an equation, in which the error rate equals 196 times the square root of $(1 / N)$, where “N” is the number of observed samples.ⁱⁱ

Number of Samples	Approximate Error Rate (%)
10	62
100	19.6
1000	6.2
10000	1.96

Of primary importance is that for monitoring traffic, the accuracy of the information is dependent on the number of samples of a given type of traffic, with accuracy improving with higher numbers of samples. The number of samples depends on the amount of monitoring time and the sample rate used to capture the packets. Users can control traffic observation to obtain the desired error rate by increasing the monitoring time and/or the sample rate to obtain sufficient samples for best results for specific environments. This capability enables a highly scalable technology, in which the sample rate can be reduced on higher speed networks without a reduction in traffic measurement accuracy.

Historically, sFlow has proven to work very well in applications that monitor network traffic levels, but has been considered unsuitable for security applications. This is because the sampled nature of the data does not lend itself to typical pattern-matching (i.e. signature-based) systems that must operate on complete packets and occasionally reassembled data streams. However, Lancope research has proven that sFlow-based anomaly detection yields significant network and security benefits.

Introduction to Flow Processing

Flow analysis for network security monitoring and host profiling is a technology invented by Lancope in early 2000. Flow analysis processes network traffic to identify suspicious behavior without the use of signatures. To that end, Lancope defined a “flow” as one or more TCP or UDP sessions between two hosts using a common service. For example, a web browser accessing a remote server will normally open several ephemeral ports to communicate with port 80 on the server. All of these sessions are aggregated into a single flow for processing purposes. Since one objective is to eliminate the use of signatures, no packet payloads are processed. Because session reassembly is not required, Lancope’s flow analysis technology offers significant performance advantages.

Initial flow processing applications used native packet captures via mirror ports and taps to gather packet headers for collections into flows for processing. The detailed information contained in the packet headers was broken down and analyzed in several dimensions to provide details about each flow sent, each host active on the network, and other activities classified as network probes or anomalies. These events were prioritized via a unique value system in StealthWatch called the Concern Index™, which is calculated based on the correlation of dozens of interacting parameters that relate directly to the characteristics of each host’s network behavior. In addition, the ability to operate on the Ethernet section of the packet provides additional details, such as the MAC address of the source and destination as well as the VLAN and MPLS tags, provide rich detail that can be valuable when tracking down miscreants on the network.

Flow analysis provides several different lenses by which flow data is analyzed. Some of these perspectives include “host centric” views that focus on all activity involving a single host; “virtual security zones” for internal and external hosts that shows traffic flows within and between local groups of hosts; and by individual flow, which identifies each flow as “normal” or “possibly malicious.” In addition, StealthWatch offers a rich set of host and zone policy parameters, which allow a network or security administrator to define and enforce policy including from bandwidth allowed, applications that can be used and “network zones” that are allowable for the user to access. This combination of multiple layers of analysis and policy provides a multidimensional network monitoring and security solution.

These overlapping, multilayered monitoring technologies enable detection of several classes of attacks, such as Denial of Service (DoS), Distributed DoS (DDoS), fragmentation attacks, various types of reconnaissance and general “user abuse” by internal users.

Flow Processing using sFlow

There are many intriguing parallels in the data provided by sFlow and native packet capture flow analysis. The most important aspect is that sFlow provides full packet headers, which contain the complete set of data processed by flow analysis. As such, processing sFlow data is an ideal fit for flow processing architectures. However, the “sampled packets” nature of sFlow can present a challenge. As described in the following sections, virtually all classes of attacks that would be detected by a flow analysis system using native packet capture would also be detected by a flow-based analysis system modified to operate with sampled data provided by sFlow. Many attacks and other forms of abuse are highly repetitive and require only a single observation to detect. Others detection algorithms for events that are by nature highly repetitive, such as DoS attacks, can be statistically modeled to prove a high degree of accuracy.

Using Sampled Data for Security

Looking at the accuracy of traffic analysis presented in the “Introduction to sFlow,” one could quickly conclude that security information available by sFlow is either error-prone or requires significant time to compile sufficient samples for meaningful data. Fortunately, neither is true. Using sampled data for security processing would be ineffective if implementation depended upon a binomial distribution as this would rarely have enough samples for any real degree of accuracy. However, different statistical rules apply when looking for security and network abuse related events. The exact analysis is highly dependent upon what is being detected as described in the following.

Protocol Violations

There are a wide range of common attacks that only require a single sample to provide 100% accuracy of detection of nefarious behavior. Examples include illegal packet fragments, fragment overflow attacks and various TCP flag attacks. If a single packet is detected with just one of these attacks, then it is likely that the packet was crafted with malicious intent. The average time to detection is directly related to the suspicious packet rate and the sample rate used by sFlow. Within StealthWatch, these types of attacks are noted both as single event alerts and also contribute to the “Concern Index” for prioritizing attacks, which alerts the operator of malicious activity and prioritizes the most suspicious hosts for immediate response.

“Dark IP Space” Access Detection

A single connection attempt to a dark space address indicates that a host has attempted unauthorized connections. Again, using sFlow, a single sample of this type of activity will result in 100% accurate detection of something suspicious or misconfigured on the network. This type of activity also contributes to StealthWatch’s Concern Index and can generate a specific event alarms to the operator.

User Policy Violations

If a flow-based detection system is implemented with an established set of user policies, then most violations of access policy by a user will be detected with 100% accuracy by a single sample. Given that users typically violate policies repeatedly, and in the case of actual abuse, then the likelihood of detection increases with higher sample rates and the amount of activity performed by the user as measured in packet counts. If this network segment is highly sensitive, Lancope recommends that a flow-based NBA system using packet capture (or packet sniffing) technology be used for the sensitive network segment to ensure that low-level activity is picked up immediately.

Peer-to-Peer (P2P) Application Detection

Detection of P2P network applications with flow-based NBA solutions has been common for several years. These technologies operate with sFlow with very little degradation as common behaviors associated with P2P applications typically include large file transmissions and connections to a large number of hosts. As such, these typically provide a sampled flow-based solution with a large number of packets can provide detection with a very high degree of accuracy.

New Host Identification

Each time a packet is sampled by sFlow, the IP address of the source and destination hosts are available. By mapping the internal IP address space, it is simple to identify when new host IP addresses become active on the network. Additionally, since sFlow provides the complete packet header information, the host MAC address is available to provide the ability to uniquely track a host IP address.

Operating System Identification

Because sFlow provides the complete packet header, passive OS identification techniques are used on TCP SYN packets to determine the probable operating system of the host sending the SYN packet. Because data flow is sampled, it takes longer with sFlow than with native packet capture to capture and identify an initial SYN packet. The delay depends on the number of new TCP connections established by the host and the sample rate used for sFlow.

Denial of Service (DoS)

DoS attacks include a broad spectrum of attacks, such as “SYN Floods,” that are typically directed at a target IP address from one or many other addresses. These attacks usually attempt to deny legitimate users access to a service by overwhelming the provider of the service. Most of these types of attacks are detected by monitoring incoming traffic and alerting when specific types of packets exceed a predetermined threshold.

With sampled data, the detection of DoS attacks is considerably more complicated. Because the data is being sampled and the detection mechanism is typically a threshold, the sample data must be scaled in relation to the sample rate used by the data source. Additionally, the calculation is based on a decaying, sliding time window to ensure that sufficient samples are available for accuracy. The resulting calculation depends on the DoS threshold for the type and target of attack, and provides a properly scaled measure of the occurrence and severity of the DoS attack.

DoS attacks also create incident specific alarms in StealthWatch and contribute to the Concern Index.

Scanning “Attacks”

Scanning is typically a reconnaissance activity that precedes an actual attack. Detection of scanning with sFlow is dependent on the scanning rate of the malware, the sFlow sample rate, and the address ranges being scanned.

The most prevalent address scanners today are worms. Worms tend to be very noisy scanners that aggressively scan a large segment of addresses, often scanning the same network segment repeatedly. Detecting worms with sFlow normally works very well, but is somewhat dependent on both the sample rate and the amount of time operators are willing to wait for detection. As a general rule, if sFlow is being used to detect and automatically mitigate worm activity inside the network infrastructure, then whenever possible the sample rate should be set to 128 or higher. This will provide the greatest sensitivity and most rapid response, allowing mitigation actions to occur quickly. However, if operators are manually mitigating via other means, a lower sample rate will suffice, but may take longer to alert on scanning activity.

Lancope recommends a sample rate of 128 for scanning detection (and hence, rapid worm detection) based on its research of internal corporate networks, which suggests that detection of scanning activity is best accomplished on class C network segments by observing connection attempts and failures over a period of time. Scanning is likely being performed if multiple connection failures occur within a single class C address block. By setting the sFlow sample rate to 128, a host that scans multiple Class C address blocks (as is common with most worms) will almost always be detected if the scanning is either sequential or random. Decreasing the sample rate will decrease the likelihood of detection during the initial scan, and may delay subsequent scans of the same address space before detecting the scanning activity of the worm. During the interim, a worm may locate another vulnerable target and propagate, whereas a higher sampling rate may detect and prevent the spread.

If this sample rate is too high for the network, then other detection mechanisms inherent in NBA flow analysis system will normally detect the worm faster than scanning detection at a lower sample rate. For example, one measure that is likely to detect the worm quickly is either the New Flows or Max Flows alarm in StealthWatch. These alarms function similarly to the DoS alarms discussed previously, but are not dependent on the flows being initiated within any particular address range. As such, the detection time is typically directly related to the speed of the worm's scanning, the Max Flow and/or New Flow alarm thresholds and the sample rate of the sFlow data source. Other techniques that include the use of detecting scanning activity in the internal dark space are also valuable aids in quickly detecting random scanning behavior with a lower sample rate.

The following screen shot from a Lancope StealthWatch Xe for sFlow appliance shows a number of hosts external to a network (in the security zone “All Outside”) with a high Concern Index (CI) attacking a local network with both SYN floods and with significant TCP address scanning:

Zone	Host	Country	CI	CI %	Bytes In	Bytes Out	UDP %	Server Profile	Client Profile	Alarms	Alerts	Operating System
All Outside	62.141.54.163	Germany	3,246,587	3,246%	0	4,600	0%			High Concern Index SYN Flood	TCP_Scan	WindowsXP/2000
All Outside	220.164.107.50	China	3,011,153	3,011%	0	0	0%			High Concern Index SYN Flood	TCP_Scan	WindowsXP/2000
All Outside	163.27.95.130	Taiwan	2,894,456	2,894%	0	0	0%			High Concern Index SYN Flood	TCP_Scan	WindowsXP/2000
All Outside	24.6.166.217	United States	1,758,918	1,758%	64,000	0	0%			High Concern Index SYN Flood	TCP_Scan	Linux
All Outside	210.151.163.13	Japan	1,643,608	1,643%	11,200	13,700	0%		smb	High Concern Index SYN Flood	TCP_Scan	WindowsXP/2000
All Outside	59.140.44.35	Japan	1,407,794	1,407%	0	0	0%		ssh	High Concern Index SYN Flood	TCP_Scan	Linux
All Outside	213.179.237.82	Ukraine	1,367,265	1,367%	0	0	0%			High Concern Index SYN Flood	TCP_Scan	WindowsXP/2000
All Outside	221.134.46.98	India	781,980	781%	0	0	0%			High Concern Index	TCP_Scan	WindowsXP/2000
All Outside	218.87.89.229	China	742,541	742%	0	0	0%			High Concern Index	TCP_Scan	WindowsXP/2000
All Outside	67.177.243.77	United States	664,663	664%	0	0	0%			High Concern Index	TCP_Scan	Linux

Since most of the scanning being done is TCP-oriented, the operating system of the attacking hosts is also available. Looking at the first host on the list, operators can drill down and examine scanning activity:

Host is Source of Potential Probes (High CI)						
Start Time	Last Time	Target IP	Hits	CI	Type(Hits)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.182.0	63	819,819	Addr_Scan-80/tcp(63)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.183.0	36	468,468	Addr_Scan-80/tcp(36)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.180.0	27	351,351	Addr_Scan-80/tcp(27)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.181.0	21	273,273	Addr_Scan-80/tcp(21)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.186.0	18	234,234	Addr_Scan-80/tcp(18)	
03/22/06 06:04:29	03/22/06 06:34:59	209.182.176.0	15	195,195	Addr_Scan-80/tcp(15)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.190.0	12	156,156	Addr_Scan-80/tcp(12)	
03/22/06 06:04:29	03/22/06 06:34:59	209.182.177.0	12	156,156	Addr_Scan-80/tcp(12)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.179.0	12	156,156	Addr_Scan-80/tcp(12)	
03/22/06 06:04:29	03/22/06 06:34:59	209.182.178.0	9	117,117	Addr_Scan-80/tcp(9)	
03/22/06 06:04:59	03/22/06 06:34:59	209.182.188.0	9	117,117	Addr_Scan-80/tcp(9)	
03/22/06 06:04:59	03/22/06 06:04:59	209.182.191.0	6	78,078	Addr_Scan-80/tcp(6)	
03/22/06 06:04:59	03/22/06 06:04:59	209.182.189.0	6	78,078	Addr_Scan-80/tcp(6)	
03/22/06 06:04:59	03/22/06 06:04:59	209.182.187.0	3	39,039	Addr_Scan-80/tcp(3)	
03/22/06 06:03:33	03/22/06 06:03:33	209.182.183.225	2	200	Reset-80/tcp(2)	
03/22/06 06:03:31	03/22/06 06:03:31	209.182.182.43	2	200	Reset-80/tcp(2)	
03/22/06 06:03:32	03/22/06 06:03:32	209.182.180.254	1	100	Reset-80/tcp(1)	
03/22/06 06:03:32	03/22/06 06:03:32	209.182.182.18	1	100	Reset-80/tcp(1)	

Based on the Potential Probes page, this host was aggressively scanning for HTTP servers on port 80 and in 30-minutes generated over 3,000,000 CI points – 30x the alarm threshold on this particular installation. The actual alarms were generated within the first couple of minutes for this host:

Most Recent Alarms				
Date Time	Alarm	Details	Ack By	Ack Time
03/22/06 06:05:00	SYN Flood	10,700 pkts		
03/22/06 06:04:30	High Concern Index	117,117 at time of alarm		

Summary of Sampled Data for Security

As shown above, the sampled nature of sFlow provides both good sensitivity and rapid detection times for many real-world events. However, for quieter events, the sampled nature of the data will tend to be less sensitive. As such, flow analysis using sFlow may be a little slower to detect as there may be insufficient samples to process. In the same vein, there is a higher possibility of a low-level event “flying under the radar.” With sampled data, the flow analysis system will see less than 1% of the total network traffic and a low percentage of all flows. Lancope research has proven that more than 90% of all flows have fewer than 32 packets. By operating with a sample rate of 1/128, StealthWatch will see less than one of every four flows. Lower sample rates will see proportionally fewer flows. As such, a short duration event can evade detection, and events, like single packet UDP worms, may take longer to detect. For an aggressive noisy event, like a worm, sFlow based flow analysis will often alarm within 2-3 seconds of a packet capture version. Quieter and slower events, that require a certain number of events to trigger an alert, may be missed by sFlow-based systems. For highly sensitive or valuable network segments where this is a concern, a system using packet capture should be deployed in conjunction with sFlow covering the remainder of the network.

Traffic Analysis

The traffic analysis performed by a flow-based NBA system will follow the accuracy table shown previously in the “Introduction to sFlow” discussion. Accuracy is highly dependent on the number of packets in the sample size. As such, larger numbers of packets are more accurate.

sFlow Sampling Strategies

As discussed, the most important considerations regarding security performance of sFlow data are the sample rate(s) and locations of the sFlow samplers. Additionally, the type of sampling being done, whether uni-directional or bi-directional, can assist in improving the data integrity. Finally, the flow analysis system can be used to monitor the traffic levels at each sFlow sampler, to provide overall network traffic analysis. These factors contribute to how sFlow can be best enabled to meet the various user requirements and will be discussed in more detail in the following sections.

sFlow sample rates are generally selectable at the port level on switches that support sFlow. In addition, sampling can be turned on and turned off on a per-port basis. This provides a great deal of flexibility in configuring the types and amounts of sFlow that will be reported.

There are a number of decisions to make prior to defining the sFlow-enablement strategy.

- Is rapid detection and automatic mitigation of scanning attacks, such as those presented by worms, a key consideration? If so, then the highest, feasible sample rate based on available network equipment and data rates is preferred. As discussed previously, a sample rate of is recommended for security applications.
- What is the maximum sFlow sample rate supported by the network infrastructure? Most Foundry equipment will support a sample rate of 128 or even higher, but some equipment is known to have lower maximum sample rates (512 or less). If not using Foundry equipment, it is best to check with the equipment manufacturer to determine the capabilities.
- Will the data be used only for security analysis or will the additional benefit of traffic analysis at the switch level be an important part of the solution? If switch level traffic is of interest for monitoring and network engineering, then it is desirable to enable sampling at each relevant switch.
- Does the network have complete sFlow visibility at all edges (ingress and egress) of the network? If yes, then full inter-host traffic and security can be performed by sampling only at the network edges with potential mitigation available at the edge port where the ill-behaved host is connected.

Following are two additional configuration options to optimize network bandwidth used by sFlow reporting:

- **Switch Statistics**
Many switches support exporting their traffic statistics on a per-port basis. Unless another application is being used to gather and process these statistics, you should turn the forwarding time off.
- **Packet Header Length in Sampled Data**
This is another sFlow configuration option in most switches. For optimum results, you should set the packet length in the captured data to 80 bytes or more, though a value as low as 60 will work in most environments. By default, many switches send 128 bytes.

These use and configuration options yield a number of different possible scenarios, three of which are exemplified below.

Case 1: All Edge Switches Are sFlow Capable

If the entire edge of the network passes through sFlow-enabled switches, then one option is to sample exclusively at the edge. This provides a complete view of all traffic passing between hosts on the network without having to enable sFlow at the network core. This case works for switches that sample uni-directionally and/or bi-directionally in any combination. The flow-based system processing sFlow should be able to detect the sampling directions and combine the samples appropriately for the environment.

With this configuration, sFlow provides a complete “security view” of the hosts on the network and the ability to map a host to the closest network switch and port. This enables mitigation at the switch port level.

By enabling sFlow on the core routers, operators can also monitor network traffic in the core, which is not a requirement for security monitoring in this configuration.

Case 2: Some Edge Switches Are sFlow Capable

If the entire edge of the network does not pass through sFlow-enabled switches, then enabling sampling at the core will provide coverage for those hosts whose sessions do not otherwise pass through a sFlow capable switch. This provides a complete view of all traffic passing between hosts on the network, and also results in some redundancy of data reported to the flow-based system. While this causes no problems as the system is capable of selecting the appropriate packets to process, it increases the number of samples that must be processed on the security appliance and creates some additional sFlow traffic on the network.

With this configuration there is a complete “security view” of the hosts on the network, and with the ability to map a host to the closest network switch and port, the capability to perform mitigation at the switch port level is available for all hosts attached to sFlow capable switches. A full set of network traffic statistics is kept for each switch and router providing sFlow data to StealthWatch Xe for sFlow.

Case 3: Only Core Switches Are sFlow Capable

If most or all of sFlow-enabled switches reside in the network core, then enabling sampling at the core will provide complete coverage for the network. This provides a complete view of all traffic passing between hosts on the network in the core which gives a complete “security view” of the hosts on the network. Mitigation at the network edge against misbehaving hosts is generally not possible with this configuration as there is insufficient information to identify and control the edge switch down to the specific port level.

Recommended Configuration

If the network infrastructure will support it, then it is best to simply enable sFlow everywhere. While there is significant redundancy in the data, this will be automatically compensated for by the flow-based analysis. By having sFlow enabled everywhere, there is a high degree of tolerance to network configuration changes that may otherwise result in unreported traffic, such as rogue or otherwise unknown network access points. If this is simply too much data for the network, the next preferred configuration is to monitor at the network edge for mitigation and inter-switch communications, and finally monitoring of the core is recommended to ensure complete network coverage – if mitigation or inter-switch communications is of no concern.

sFlow Bandwidth

Given that each sFlow packet contains up to 11 sFlow samples, the increased network load is very light compared to the monitored traffic. With a sample rate of 1/128, sFlow sends one packet of up to 1500 bytes for each 1408 packets observed with a header length of 60. Given a nominal packet size on the network of 500 bytes, this will increase the network traffic by 1500 bytes for each 704,000 bytes on the network – about 0.21%. That is, monitoring a fully loaded 1Gbps link with a sample rate of 1/128 will generate about 2.1Mbps of sFlow traffic. When working with a large number of slow sFlow devices, the overhead of sFlow will have an additional impact. As such, for planning purposes, a value of 0.5% is a safe number to use; this would translate to 5 Mbps of sFlow traffic for each one Gbps of monitored traffic.

Monitoring High Speed Network Links

This paper has focused primarily on the desirability of using higher sampling rates to improve detection capabilities within corporate networks. In general, a higher sample rate is best, particularly when monitoring internal hosts for policy violations. However, use of sFlow to monitor high-speed trunks, where a lower sample rate may be required due to equipment limitations, provides some significant capabilities that may meet or exceed the customer requirements.

Sample rates of less than one packet in 1024 will provide useful data for traffic reporting and for DoS/DDoS detection. As discussed in the “Introduction to sFlow,” the accuracy of traffic information is dependent only on the number of observed samples of a particular type. Therefore, the traffic accuracy of a system that has sampled 1000 packets at a sample rate of 128 will be the same as a system that has sampled 1000 packets at a sample rate of 8192. The key parameter to control is the amount of time required to accumulate sufficient samples for ideal accuracy.

DoS and DDoS attacks on the network, such as high speed trunk servicing a server cluster sharing a common IP address, may be detected as well for configurations. By monitoring the number of received SYN packets directed at the cluster IP address, both DoS and DDoS attacks will be detected with a high degree of accuracy. Note the default SYN flood threshold provided by StealthWatch should be adjusted to a higher value that matches the infrastructure.

StealthWatch System

The StealthWatch System, which includes StealthWatch NC for native flow capture, StealthWatch Xe for NetFlow, StealthWatch Xe for sFlow and StealthWatch Management Console, is the leading Network Behavior Analysis (NBA) and Response that defends internal networks against zero-day attacks, internal misuse and unnecessary network exposures. StealthWatch continuously monitors network behavior, detects anomalies, and isolates known and unknown threats. Expanding beyond the capabilities of traditional network security products, StealthWatch collects, categorizes and analyzes network traffic to create comprehensive security intelligence at both network and host levels. Providing a cost-effective, single point of reference for optimizing security and network operations, StealthWatch enables organizations to improve the network health and security posture of their networks.

StealthWatch for sFlow: Competitive Differentiators

StealthWatch is currently the only native sFlow Network Behavior Analysis (NBA) solution available. Some vendors in the NBA space offer a rudimentary sFlow capability, which is actually done by taking the sFlow samples, converting them to a NetFlow record, and discarding all of the packet level information that is provided. Other software providers focus on the traffic analysis available from sFlow to provide some rudimentary security knowledge via connection analysis to detect worms and a very limited signature capability that works only when the correct packet is sampled. Following discusses the advantages that StealthWatch provides over both types of systems.

NBA Solutions using sFlow

There are a number of advantages provided by StealthWatch using “native sFlow” as opposed to other solutions that convert sFlow into a NetFlow data stream:

- *Packet Level Information*
Processing native sFlow data preserves all of the packet header information, including host MAC addresses (which are meaningful for edge switches), TCP header flags and option fields, VLAN and MPLS tags, TTL values and some of the packet payload. In addition to the inherent advantage of having the packet header information, this can also be processed to identify the host operating system, whereas converting sFlow into NetFlow loses this information.
- *Statistical Significance*
By operating on native sFlow data, StealthWatch can process the information both with and without scaling, which allows for more granular and detailed processing and the ability to easily and automatically deal with multiple sample rates. When sFlow is converted to NetFlow, the data is either scaled by the sample rate or it is not. If scaled, then the ability to process unscaled data is lost. If not scaled, then the operator has to inform the processor of the sample rate, which is subject to change and error.
- *Flow Processing*
Creation of flows for additional processing is much more consistent with native sFlow than conversions. Similar to native packet capture, sampled packets are accumulated into flows when processing native sFlow. This assists both in providing an improved picture of client/server interactions, as well as improving the usefulness of the forensics data provided by sFlow. In contrast, when the sFlow samples are converted into NetFlow, each individual sampled packet becomes a NetFlow report, which in turn must undergo significant additional processing to be rejoined with other packets from the same flow to create an accurate flow report.
- *Data Reporting*
sFlow also provides a number of enhanced data reporting capabilities that provide a roadmap for future enhancements. For example, “802.1x username identification” is supported by sFlow. However, no facilities exist to convert this type of information into NetFlow.

Conclusion

A few well-known products that use sFlow for traffic analysis have attempted to move into the security space. In general, these efforts have fallen into two general categories. First, some limited signature capability can be used to identify worms and other very noisy, well-known events. These attempts are limited in that the correct packet must be sampled, the signature for the known event must fall within the fraction of the packet that is sampled, and, of course, the signature must exist. The second category has been to do some rudimentary connection analysis to identify likely hosts that may be infected with a worm or performing a DoS attack. These approaches pale when compared to a full-blown StealthWatch System. The primary benefit of deploying StealthWatch is the ability to fully understand all activity performed by a host, including ports used, bandwidth usage and the full range of malicious activities StealthWatch is designed to monitor. StealthWatch then maintains a complete forensics log that allows event investigation and resolution – essentially providing the contextual information necessary for a complete investigation. Furthermore, StealthWatch provides a complete reaction infrastructure is in place that allows alerting and response activities to these events in a manual or automated fashion.

StealthWatch Xe for sFlow provides a unique and powerful security solution for the market segment that relies on sFlow for its network management and reporting infrastructure. Significant advantages exist in the approach developed by Lancope that make StealthWatch the only reasonable choice for Network Behavior Analysis and Response in these environments.

About Foundry[®] Networks

Foundry Networks, Inc. (Nasdaq: FDRY) is a leading provider of high-performance enterprise and service provider switching, routing and Web traffic management solutions including Layer 2/3 LAN switches, Layer 3 backbone switches, Layer 4 - 7 application switches, wireless LAN and access points, access routers and metro routers. Foundry's 8,500 customers include the world's premier ISPs, metro service providers, and enterprises including e-commerce sites, universities, entertainment, health and wellness, government, financial, and manufacturing companies. For more information about the company and its products, call 1.888.TURBOLAN or visit www.foundrynetworks.com.

About Lancope[®]

Lancope is the leading provider of network behavior analysis (NBA) and response solutions that defeat zero-day worms, internal network misuse and other anomalies that compromise network integrity. Lancope's StealthWatch System integrates security and network management technology to reduce network risks and maximize network availability by rapidly identifying, prioritizing and mitigating critical threats, whether new or well-known. Both OPSEC and Common Criteria-certified, StealthWatch was named an InfoWorld 2005 Technology of the Year. Defending the networks of Global 2000 organizations, academic institutions and government entities, StealthWatch protects over 200 enterprise customers, more than all direct competitors combined. Lancope's Technology Alliance Partners include Foundry Networks, ArcSight, IBM Tivoli, LURHQ and CheckPoint. Lancope is a privately held, venture-backed company headquartered in Atlanta, Georgia. For more information, call 888-419-1462 or visit www.lancope.com.

About the Author

As co-founder of Lancope and Chief Technology Officer, John Jerrim is responsible for development of the real-time, flow-based engine behind StealthWatch, the award-winning intrusion detection system (IDS). Prior to co-founding Lancope, Jerrim operated his own communications software company for ten years, where he consulted and developed communication and diagnostic software products.

Previously, Jerrim was Systems Engineering Manager and managed PC software development at Hayes Microcomputer Products, where he was responsible for a line of communications software products as well as the development of advanced PC communications hardware. He began his career at Sangamo Weston/Schlumberger, where he advanced to Manager of the Advanced Technology Laboratory, North America.

Jerrim holds seven US patents and has several pending in communications, network security and real-time processing. He has a BS and a MS in electrical engineering from Clemson University and is pursuing a Ph.D. in the same field from the Georgia Institute of Technology.

i "History of Packet Sampling", http://www.sflow.org/about/sampling_history.php

ii "Packet Sampling Basics", Peter Phaal and Sonia Panchen, sFlow.org, 2002, <http://www.sflow.org/packetSamplingBasics/index.htm>