plixer International

Flow-based Approaches in Network Management

Authored Jointly by Plixer International and Cisco Systems

Whitepaper

Executive Summary

The continuous increase of gigabit speeds in computer networks has considerably stimulated the usage of flow monitoring techniques for network management. For this reason, researchers and operators are searching out more flexible and scalable solutions.

Insight and awareness are the primary drivers pushing flow developers to serve up deeper details that some believe are teetering on the edge of leading the IPFIX standard to emerge as packet capture. Where will the proverbial line in the sand be drawn that separates packet analysis from flow analysis? Where is flow technology going? How is it evolving and what enhancements are readily available today?

The largest market for flow technologies has and will continue to be investigation. This document explores what this means to threat forensics and application awareness. It discusses how filtering and reporting go hand and hand and why combining this with comprehensive archiving is doing an effective job at alleviating many pressures related to regulatory compliance.

Collecting ever increasing metrics at gigabit speeds means massive volumes of flows. This study outlines distributed collection as well as alternatives to flow and packet sampling. It also provides details on how the unique needs of both the cloud computing and virtualization industries are being addressed. For service providers who view Software Defined Networks as on the horizon, this paper explains how NetFlow v9 and IPFIX will continue to play a key role here as well.

The goal of this document is to highlight what is available in the market: here and now for network management and monitoring. It provides specific examples of where is the innovation seems to be headed and the challenges that lay before it.

Content

Why the Explosion in NetFlow and IPFIX Growth

Insight is the driver. With a consumer base thirsting for details on who, what, when and where, NetFlow and IPFIX are the network traffic analysis technologies that meet nearly all of the consumer demands. Why has the market selected flows over other technologies such as SNMP or packet capture?

Wide Adoption

Flow data like SNMP is widely adopted and has been implemented on nearly every router manufactured by most vendors. This is because it has several advantages over traditional traffic analysis protocols. Flows are a push technology whereby, unlike SNMP no polling is needed. For this reason, it also doesn't require a password or community string. It also isn't encrypted unless SCTP is used over UDP as the transport.

IPFIX, NetStream, J-Flow, AppFlow, Cascade Flow, etc. are all nearly exact copies of NetFlow and referenced collectively as flow technologies in this document. IPFIX however, is the official IETF standard for all flow technologies.

Accessible

Accessibility is another driver of flow growth. It's readily accessible in almost all the corners of every network. It simply needs to be turned on. Nearly all firewalls support it as do many server operating systems (e.g. VMware, Linux). Several vendors are now marketing true flow capable switches in lieu of packet sampling technologies such as sFlow.

Inexpensive

It's inexpensive. In most cases, flow technology is integral to the devices and simply needs to be turned on. In contrast to expensive packet analyzers that have to be purchased,

deployed and maintained, flow technologies are free, part of existing maintenance programs and provide insight into more areas of the network. According to the Gartner Group flow technologies should be done 80% of the time and packet capture with probes should be done 20% of the time[1].

Details

When most consumers think about flow data, the details available in NetFlow v5 come to mind. The technology however, has come a long way from the 20 or so elements in NetFlow v5 to the tens of thousands available in NetFlow v9 and IPFIX. Flow technologies are now being used to export such details as: system messages, CPU utilization, round trip time, HTTP Host, URLs, packet loss, retransmits, jitter, VoIP codec, caller ID, layer 7 application, TCP window size and much more. This is a technology that is starting to rival the details previously only available through packet capture. TWith the advances since the Gaernet report was issued in 2012, today flow analysis can be used 90% of the time and packet analysis only 10% of the time. This is not to say that flows will completely replace packet analysis because it probably never will.

As the elements of what flow technology can be used to export proliferates, it carries an Achilles heel: volume. The more details administrators try to stuff into a single flow tuple, the more overhead is produced. For example, requesting URL could cause what was a single flow to be broken up into multiple flows. To add salt to the wound, more details means more bytes pushed into the same flow. A single NetFlow datagram used to send 30 flows could accomadate only 4 flows if excessive details are quested from each flow. Asking NetFlow v9 or IPFIX to export greater details starts encroaching on packet capture turf and begins to defeat one of the underlying intentions of flow technologies (i.e. less is sometimes more). For this reason, the trend in the industry that is supported by multiple

1 http://www.gartner.com/id=1971021

hardware vendors is to allow the user to select what they want to export.

The third area that needs consideration when expanding the flow tuple to include more details is overhead. Asking the devices (e.g. routers) to match on more criteria and provide more information about each flow places greater overhead on the hardware. When the tuple is fixed, the processing can be done in ASICs however, if the vendor chooses to make the flow fields (i.e. elements) definable by the end user, it generally requires much more CPU.

The largest market: Investigation

What are consumers doing with NetFlow and IPFIX? Historically, it was largely being used reactively to follow up on complaints related to network performance. In the past few years it is being used more and more to follow up on potential malware activities. For these reasons, flow technology has had to evolve in order to both help IT determine whether slowness is attributed to utilization or response time as well as assist with forensic sleuthing. To better deal with these objectives, NetFlow v9 and IPFIX introduced new metrics not available with NetFlow v5.

Perhaps the most popular new details being exported in flows today are layer 7 application name (e.g. Skype, Facebook, Citrix, WebEx, etc.), round trip time, retransmits and HTTP Host/URL. Although not all vendors support these elements, popularity is growing. This exciting new paradigm has also introduced a problem because many collector vendors need time to update their decode engines to support the new elements. Although IPFIX does make provisions for dynamic element adoption in RFC 5610 only a couple of vendors have added it to their export. Without support for RFC 5610, IPFIX faces the same conundrum that that packet capture and SNMP suffers from: the need for decode libraries and MIB files respectively.

What makes good reporting

Displaying the flow data is one of the many criteria that set flow consuming vendors apart. Reporting needs to provide the ability to narrow in on problems fast. For this reason, good flow reporting is generally defined by two important traits: filtering and report options.

Filtering

Narrowing in on the desired traffic is sometimes a process of trial and error. During the process, filters are added that include or exclude certain attributes. If the desired results are not the outcome, the filter is removed and another is added. Most vendors provide the user with a fixed number of default filtering options ranging that include: IP address (es), ports, autonomous system, interface, next hop, etc. The choices are pretty much limited to the information exported in NetFlow v5 which can be very limiting.

A better strategy taken by some reporting vendors is to allow the consumer to leverage any element that is exported in the flow template. For example, if MAC address, packet loss, round trip time, URL, etc. are in the template, the user can add the element to the filter, match on specified criteria and include the value based on logic such as: greater than, less than, equal to, not equal to, like, etc. Usually, numerous filters can be added to the report in an effort to focus in on the desired flows. Filtering alone however, isn't enough.

Report Options

During the process of adding filters, often times the report needs to change in order to expose information that the user needs to filter .on. At times, users want to begin the investigation process by removing 'noise' around the selected time frame. By reporting on ports or protocols, the administrator can exclude subnets or protocols such as TCP and ICMP in order to focus on UDP or other traffic type.

Some reporting solutions provide a custom report creator which empowers the user to choose the elements from a selected template as well as specify the columns and the values to group by in the report. These are the tell-tale signs of an advanced flow reporting and analysis solution as going back to a vendor with a new report request may be met with a quote for professional services. Below is an example of a solution that allows the user to filter on any element in a selected template (e.g. HTTP Host).

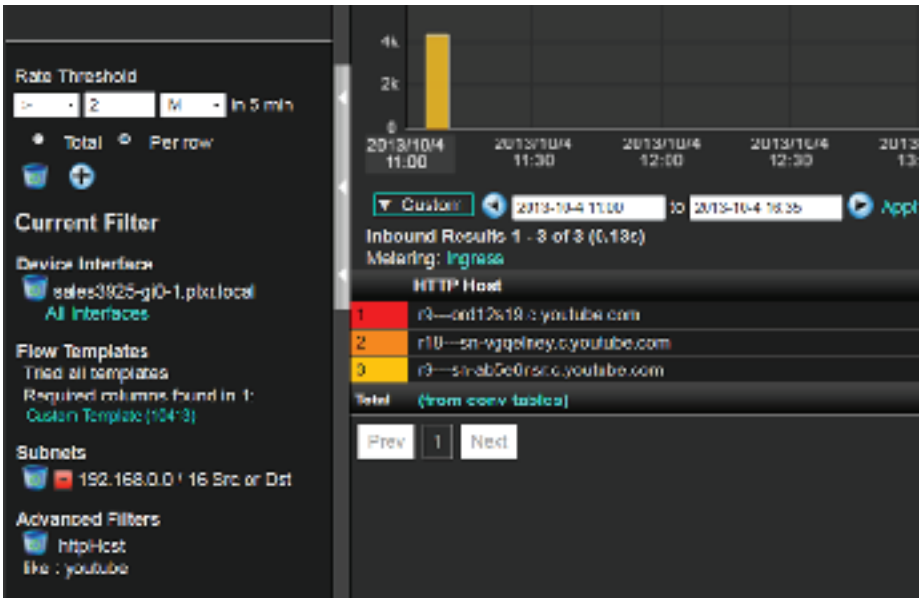

Investigating threats

Beyond determining why or what is causing slow network connections, flow technologies are used for sleuthing down suspicious activities that might become network threats. A reporting solution's ability to allow the user to navigate down to exactly what they are trying to uncover is based on two primary factors:

1) The abilities of the end user and his/her familiarity with the tool;

2) The reporting and filtering capabilities of the in-house solution.

As stated earlier, packet probes are generally not readily available in all areas of the network. Only NetFlow and IPFIX provide the all-encompassing view into every corner of the infrastructure. When end to end visibility is paramount to understanding how malware

entered and moved laterally within the organization, every router and switch becomes a probe. In a sense, we can think of every flow capable device as a video camera that streams captured footage back to the DVR (NetFlow collector). In many security fields, the cameras are the first place people look for evidence surrounding a questionable event .

Regulatory Compliance

Concerns over regulatory compliance such as HIPAA, FIPS, NERC, SCADA, PCI, NPPI, SOX and COSO have executives prepared to make investments to ensure that in the event of an audit, they can provide the deepest levels of visibility, accountability and measurability required for ensuring and maintaining compliance with these industry standards. In order to avoid hefty fines, flow data can provide the details necessary to:

- Identify connections to and from the SCADA network
- Track and account for healthcare employee network activity
- Recognize unauthorized host access enabling rapid response for electronic protected health information (EPHI) access, alteration and/or destruction
- Detect malicious and suspicious network activity
- Leverage third party integrations for threat mitigation to remediate security policy violations
- Profile hosts for violations of security policies
- Continuously monitor hosts and network activity to identify intrusions
- Ensure and optimize SCADA network and application performance, availability and internal security
- Leverage user accountability for security and network risk visibility
- Measure and prioritize risks
- Conduct forensic analysis for security incidents

Flow technology allows administrators to quickly confirm the source of the problem by narrowing down

the issue to the client, server or network.

Organizations which impose 'locking' policies that define which groups can and cannot communicate with one another can leverage flow data to verify that there are no communications between prohibited groups of users/hosts. If rules are violated, an alarm is raised and full audits can be run to report on all end systems involved. Given ample disk space, some flow solutions can save all raw flows from all flow exporting devices for decades. This ability is a crucial component of an organization's compliance documentation.

Virtualization, the Cloud and SDN

These growing and potential markets are well supported by flow technologies. By doing almost nothing specific to these industries and staying focused on the needs of end users, flow technology is in most instances the protocol of choice for managing these new areas of communication. The same flow details already referenced in this document are applicable to these emerging markets.

Virtualization

According to the Gartner Group, the adoption rate of server virtualization will reach 82.4% of total OSs in 2016[2]. This phenomenal growth rate has been great for early adopters of flow technologies for a couple reasons:

- Each ESX license natively has the ability to export IPFIX
- Flows are the only way to monitor some applications installed on VMware

Due to technologies such as "Distributed Resource Manager" in VMWare, the physical location of an application as well as the servers it resides on can change dynamically. This auto movement of the application can make monitoring the traffic in and out of the application difficult especially if the routers bordering the application are no longer passing

2 http://www.gartner.com/id=2210316

the traffic. By enabling flow technology on the servers that are supporting the application administrators can still obtain a full view of the entire application. There is of course a downside.

Virtual OS developers like VMware have been slow to adopt the innovations that the industry has seen from other vendors. Although they support IPFIX, it is largely limited to the same 20 or so details exported in NetFlow v5. If consumers are migrating to VMware in order to save money and to possibly obtain performance enhancements, metrics on round trip time, packet loss and retransmits ideally should be included in the mix of elements that an ESX server could export. It is expected that the market will put increased pressure on virtual OS companies to export richer contextual details on the applications and users connecting to the servers they develop. A good flow solution should be ready to take advantage of it.

Cloud Computing

Cloud computing is exploding. Gartner stated that the public cloud services market is forecast to grow 18.5 percent in 2013[3] and to continue growing to $3.1 billion by 2015[4]. With more companies outsourcing to the cloud, how will they monitor and verify that acceptable service levels are being met? The same flow technologies mentioned earlier are providing assistance here as well.

Once management has ascertained the necessary availability and response times needed to optimally service the business, NetFlow and IPFIX can be utilized by a flow analyzer to set thresholds at acceptable service levels (e.g. latency). If excessive degradation of individual or collective (i.e. subnet) cloud service connections are witnessed by the analyzer, notifications can be triggered. There is however, a looming issue that is increasing the difficulty to monitor cloud application service levels. It's called SSL.

3 http://www.gartner.com/newsroom/id/2352816

4 http://talkincloud.com/cloud-computing-research/gartner-cloud-security-market-grow-31b-2015

Software Defined Networking

NetFlow and IPFIX will continue to play an important role in Software Defined Networking (SDN) because they provide accountability and network traffic visibility. Perhaps the biggest opportunity in SDNs is in what some documents refer to as the controller. The controller is the server which maintains the policies on how traffic is passed in the network fabric. If the switches and routers participating in an SDN send information back to the controller regarding the traffic it is making decisions on, the SDN controller may compile this information and export it as IPFIX. Controller exports may include details with statistics on how applied rules are performing globally. Additionally, if flow exporters tag flows with the policy used to make a decision on individual connections, reporting solutions could provide more insightful performance details. It's time to wait and see….

What's the bottom line on SDNs? The introduction of SDNs onto the network will occur at a pace similar to what many observed with Linux and virtual servers. In other words, generally new technologies are deployed out of a business need. To deploy an SDN, businesses must ask themselves: is traditional networking not meeting the needs of the business applications? Most would say it is but, the keepers of some of the largest cloud services believe they have a need for SDNs today. Either way, how flow data is exported, collected and displayed won't change much.

Layer 7: Application Awareness

In the early years of NetFlow, applications were identified by looking at the source and destination port of a connection. Many NetFlow analyzers would apply the following logic when looking at the ports:

- Which is lower, the source or destination port[5]?
- Is the lower port defined in the IANA defined port database ?
    - If yes, give it the name found

5 http://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers

o If no, look up the higher port. Is it defined?

o If yes, give it the name found

o If no, simply display the lower port

For example, if a flow was identified as using TCP source port 5003 and destination port 80 and 80 was defined in the database, the application for the flow is identified as HTTP. If however, 80 wasn't defined and 5003 was, the application would be called FileMaker. If however, neither port was defined, then the application would be labeled as TCP 80. In the early years of NetFlow, this strategy enjoyed a decent degree of accuracy. Today however, TCP 80 has become the transport of choice for most Internet destined applications! Because of this paradigm shift, the industry needed a more reliable approach to application identification. Along came DPI.

Some routers and firewalls perform Deep Packet Inspection (DPI) in an effort to try and

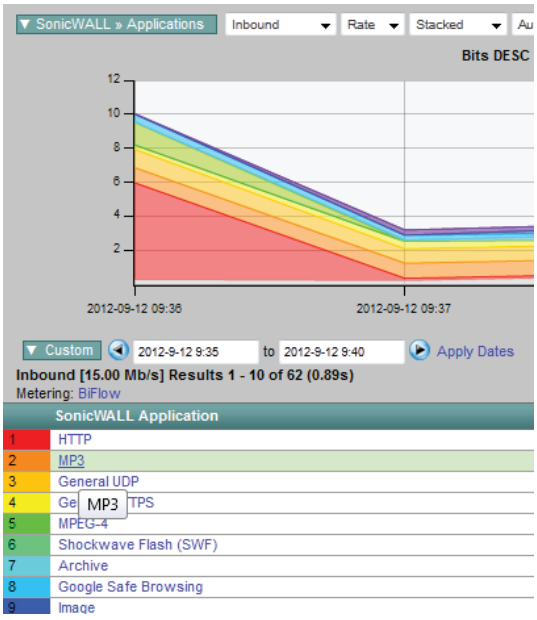

more accurately identify different layer 7 applications (e.g. Flash, Google, BitTorrent). Often times these applications share the same ports. DPI accomplishes the identification by observing a series of packets within a flow and matching the behavior to an application profile. See below:

The problem with DPI is that most vendors fail to consider SSL encryption. As a result,

applications hiding behind an HTTPS connection can evade identification by most DPI engines. One vendor has taken the time to write an SSL DPI capability which can identify applications such as Facebook.com that are hiding behind SSL. Hopefully other vendors will follow suit.

Network Threat Detection

Although flows don't contain the entire packet, they do contain the details necessary to detect many types of threats such as network scanning, receiving ICMP redirects, participating in a denial of service attack and dozens of other unwanted behaviors. Some vendors even use flow data to profile normal behaviors on the network then in turn, when a host communicates outside of its behavior in the past, an event is triggered. In more sophisticated threat detection systems, events don't trigger notifications directly; rather, they increase threat indexes which can lead to notification.

Single Event Notification and False Positives

If a single device on the local network reaches out to the Internet to a host with a reputation of being part of a botnet, this one event can but, usually does not mean that the machine is infected. If the same device also receives a few ICMP redirects from the router supporting the subnet, a security admin still can't discern that an infection exists but, suspicions would be rising. Notification based on any one event often leads to false positive notifications.

Look for a Series of Events

If a machine is routinely reaching out to known bots on the Internet and starts scanning ranges of IP addresses, scanning specific hosts or starts communicating in ways that are not typical of its normal behavior, the admin still can't definitively deduce that the device is hosting malware but, it may be time to take action and look more closely at the suspected device. In the world of threat detection with flows, reacting to any single odd behavior

generally leads to tail chasing because in data networking, normal communications often lead to an occasional odd connection.

Given enough data and time, some security appliance somewhere will ascertain that the host is distrustful. When this happens that threat detection solution may trigger what's commonly known as a false positive. If a solution serves up excessive false positives, it becomes associated with the boy who cried 'wolf' and when that happens, admins perceive little to no value in the appliance when trying to positively identify malware. No one wants to chase their tail.

Associating infection with a single event is an effective means to identifying malware just as a blood test can positively identify many viruses. But, threat detection can't rely solely on single events to stop all insurgencies. To be more effective at keeping the network clean of sophisticated intrusions such as Advanced Persistent Threats (APT), security administrators must consider the collective odd behavior episodes from every machine on the network. This is done through the use of threat indexes.

Threat Indexes

The idea behind threat indexes is that they rise for an individual host each time it participates in a behavior that is suspicious. Depending on the type of behavior (e.g. scanning the network) the event may increase the index by a higher value than others (e.g. receiving an ICMP redirect). If the threat index of a host hits a threshold, a notification can be triggered. Keep in mind that the index is a moving value because individual events age out over time. For this reason, a host with a threat index must reach the threshold within a configurable window of say 14 days because the same events that increased the counter are also aging out and as a result, the index could be reduced.

What's key to the threat index approach is all encompassing awareness. To obtain this, administrators need flow data and ideally - all of it. Threat detection at this level can't be limited to what is coming and going to the Internet. It needs to be fed all internal communications from all networked devices to ensure that all hosts - not just laptops, servers and BYOD - are being profiled.

Scaling NetFlow Collection

As exported flows become more detailed with richer metrics, flow volumes will likely increase dramatically. To deal with the increase in flow volumes, most vendors suggest deploying multiple collection points. Obtaining complete visibility across all collectors from a single interface introduces the concepts of flow deduplication and flow stitching across collectors.

Deduplication

NetFlow Deduplication is performed when multiple routers/switches export the same flows for two hosts communicating with one another. If a flow (e.g. A to B) with a matching tuple is collected by 3 routers, the collector will save only one instance of the flow with a few details.

During flow deduplication, some of the details are either dropped or assumed (e.g. DSCP values, packet and byte counts, next hop, etc.). Because the process of deduplication results in lost details, the original flows are often saved as well. Hence, the belief that flow deduplication saves disk space is false unless of course network admins are willing to work without much of the juicy details that make NetFlow and IPFIX attractive. In truth, NetFlow deduplication often leads to greater disk space consumption; however, it does play a very important role in scaling flow collection.

NetFlow Stitching

NetFlow Stitching is the process of looking at certain protocols (e.g. TCP) and assuming that a return flow occurred. If the return flow is found (E.g. B back to A), the details are merged into the A to B flow making it a single bidirectional flow. This of course is dependent on "If the return flow is found" because the return path could be going through another router which is exporting flows to a different flow collector. In most multi-collector environments, flows aren't stitched across collectors. For this reason, it is best to perform stitching and deduplication during the ad-hoc query functions.

Sampling

Most network administrators don't like the idea of sampling but, agree that at some point it seems inevitable. Modifying the flow tuple used to match on packets is an excellent strategy that allows admins to avoid sampling by compromising on only a couple of elements (E.g. removing source and destination ports from the export). If application visibility is imperative, adding a single element back called layer 7 application will provide the majority of details required without reintroducing the excessive flows problem. In some cases a 90% reduction in flow volumes can be obtained with this strategy while still maintaining 100% accuracy.

Innovations in Flow Technologies

Where is flow technology going? To understand where IPFIX is being taken, it is important to appreciate how its predecessor NetFlow v9 evolved. NetFlow v9 introduced the concept of a template which allowed the hardware to export just about any information desired. Examples of this included:

- message logs from firewalls
- meta data with details related to flows such as layer 7 application

- option templates with details on:
    - o the statistics of flows being exported
    - o the hardware interface names of the router/switch
- it can export entire datagrams like sFlow

In short, NetFlow v9 paved the way for IPFIX.

Initially IPFIX was a near copy of NetFlow v9 with several improvements, two of which have already played a major role in the industry.

- The ability to support different vendors through the use of the enterprise ID. NetFlow is 100% owned by Cisco. Any effort by a vendor to specify new elements could get trumped by Cisco.
- The ability to support variable length strings for items such as message logs and URLs

NOTE: Cisco did carve out space in NetFlow for several vendors and ultimately gave NetFlow to the Internet community which led to IPFIX.

New IPFIX Exports

IPFIX has been leveraged in several exciting new areas of the industry. Below are a few examples of how IPFIX is being used to export information:

- Some devices which historically exported system messages as syslogs or SNMP traps are now exporting these details in IPFIX.
- Flow Replicators have been introduced which allow a single stream of NetFlow or any other type of UDP data to be transparently replicated to multiple destinations. The source IP address never changes. Only the destination IP address is changed to the target host. Therefore, the receiving host believes it received the data gram directly from the ingress device. A Flow Replicator can forward the same flows to multiple destinations, so it helps companies with regulatory compliance by ensuring a backup of all system messages and notifications should an audit become necessary.

- IPFIX Gateways can act as a syslog to IPFIX converter by listening for syslogs, extracting the details and then forwarding them on inside structured IPFIX datagrams.
- The IETF has engineering efforts underway to export SNMP details in IPFIX datagrams.

New IPFIX Presentations

The uses found for new IPFIX exports continue to grow. Because the contents of IPFIX can convey far more than traditional NetFlow v5 contents, the way the data is displayed has also had to evolve. Below are examples of how flow data containing unique exports has been displayed.

- Matrix: This view provides a unique perspective on the volume of hosts communicating with a single entity. The matrix can be turned, zoomed in and hosts can be clicked on to see inbound (green) and outbound (blue) connections.

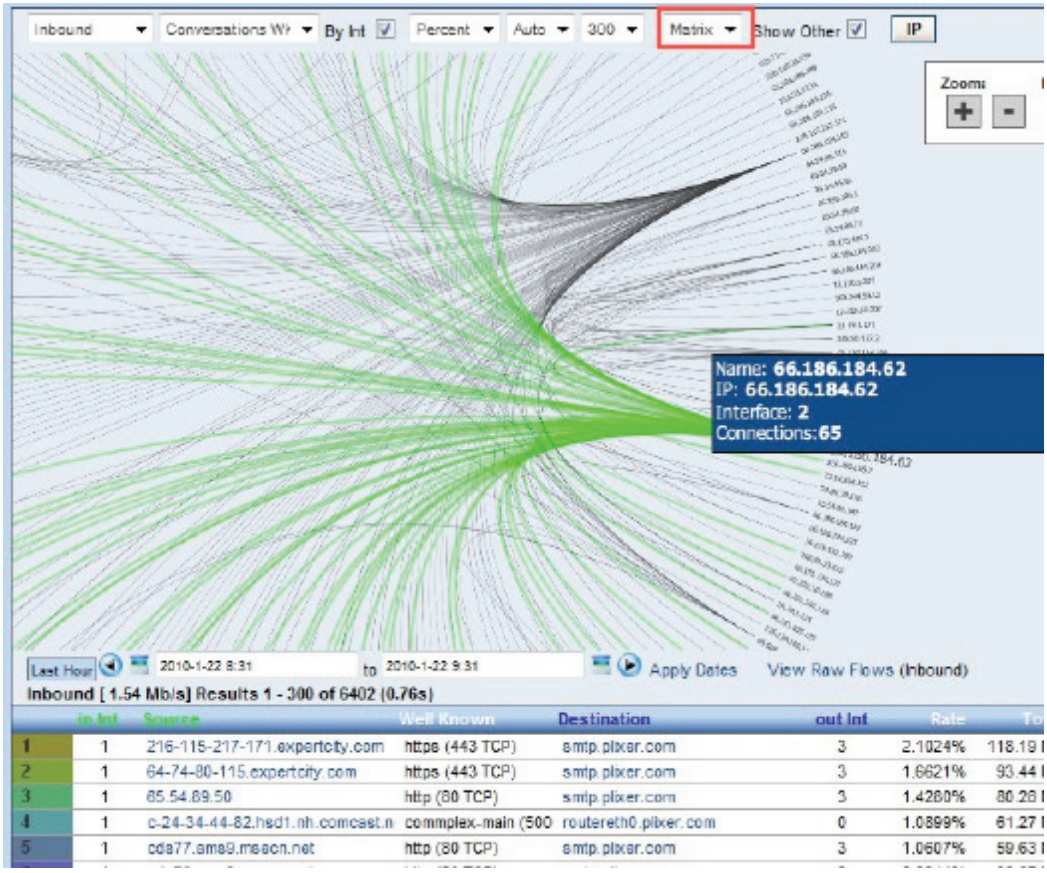

• Sonar: The threat sonar puts the geo location of the local browsers Internet IP address in the middle and pivots a radius arm around in a circle which highlights hosts found in flows with poor Internet reputations.

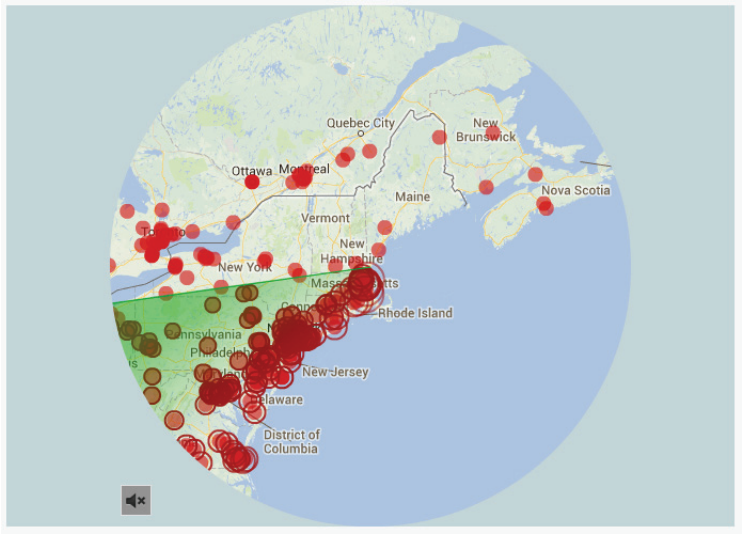

• Flow Path: These solutions display the entire hop-by-hop path of a flow. One solution works with nearly all versions of NetFlow, IPFIX and other flow technologies. Users can click on each hop along the way to see details on how the flow may have changed. Root cause analysis becomes easier when changes in DSCP, TTL or even packet loss and jitter are simple to identify.

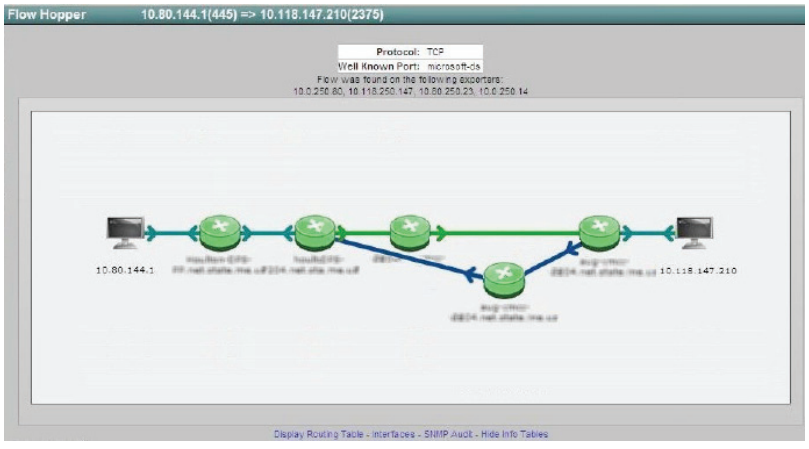

• Topology: Displaying the utilization of links graphically is not new in the industry. Displaying the overall latency or performance of a specific application is. One vendor does it all with flow data and animates the links.

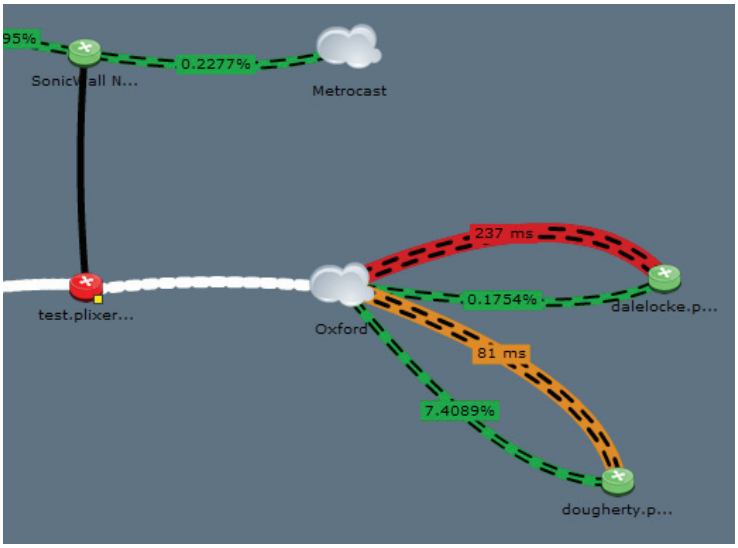

• Heat Map: One strategy for detecting malware is a process where flows are passed through a series of algorithms which look for odd communication patterns. Suspicious connections are identified by looking at behaviors in the flow ratios related to TCP flags, port volumes, comparing activities to stored baselines and other proprietary forensics. The algorithms violated carry different weights and will increase an individual host's Threat Index (TI) shown below. If the index reaches a threshold prior to events aging out, notifications are triggered. Hosts displayed high and to the right are the most suspicious.

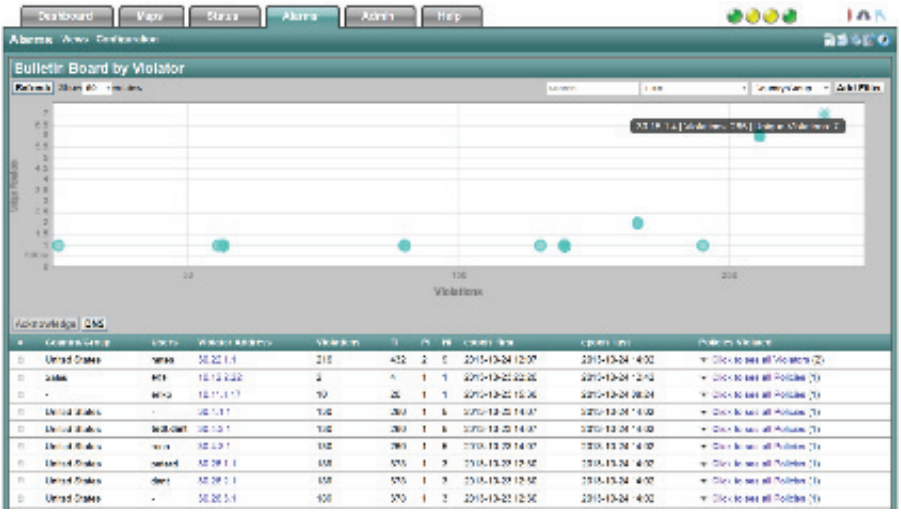

Summary

NetFlow and the newly ratified IPFIX IETF standard have become recognized as flexible and readily available protocols for improving several network management related tasks. Flow technology has come of age and is no longer viewed as a simple accounting protocol for billing or for just identifying bandwidth hogs. The introduction of templates has opened up the technology and empowers hardware developers to export nearly any metric / information imaginable.

Due to its widespread adoption and surge in innovative exports, flow technology can be used to help optimize applications at nearly all layers of the OSI model. When it isn't employed by hardware for monitoring performance, flows can carry details related to syslogs, event logs and nearly any other machine message. Once centrally located, messages can be correlated across different vendor platforms and archived indefinitely for historical research which also services several regulatory compliance obligations.

Two of the largest opportunities facing IPFIX analysis are in the areas of distributed collection and threat detection. Like many technologies that see continual improvement, expect flow volumes to grow as consumers demand more insight. How to collect and centralize the data will become increasingly difficult.

Because flows don't contain the actual packets, behavior monitoring over a series of events poses the largest threat detection opportunity with flow data. Distributed collection and localized threat detection will require centralization of triggered events. Vendors who recognize these trends are already working toward solutions that address all of these issues.