
ixia
The
ABCs
of Network
Visibility

ABCs of Network Visibility

Table of Contents

Introduction

This book is a reference guide for individuals who want to learn about network visibility and monitoring. It is a compilation of blogs from the Ixia website written by various subject matter experts in the area of network monitoring and visibility.

The definitions contained within provide a starting point to understand network monitoring and network visibility. This includes a common set of components that are part of a network visibility architecture. While some of the terms can be intimidating at first, hopefully these straightforward definitions will make those terms easier to understand.

Definitions are arranged alphabetically for easy lookup. A simple "Definition, Use Case, Considerations" format allows you to quickly understand what the item is, how to use it, and then provides some items to consider when making a purchase for this type of functionality.

For the purposes of this, some basic background information is required. This book discusses various visibility topics. Most people think of visibility as the security and monitoring tools that perform application performance monitoring, or network performance monitoring, or big data analytics. It is more fundamental than any of that. Network visibility is how data is collected, aggregated, distributed and served to those monitoring and analytics tools. Network visibility is also about creating a stable foundation for your security infrastructure.

For the content of this book, we will use the following visibility architecture as a reference. There are three basic components to a visibility architecture – the access layer, control layer, and the traditional monitoring tool layer.

The first layer of the model is monitoring data access layer. This is provided by some sort of network access device, whether it's a physical tap, virtual tap, or mirrored port. The next layer uses a network packet broker, also called an NPB, which allows you to filter, aggregate and load balance data. From a business point of view, there are several high level benefits that packet brokers provide. Which include: connectivity, reduced tools costs, scaling, reliability, and longevity. After the packet broker are the monitoring tools. This is probably what most people are familiar with. Instead of receiving the data directly from a tap, these tools now receive filtered data from the NPB that is more relevant and concise. NPBs are designed to augment your monitoring tools, not replace them, to make the tools more efficient.

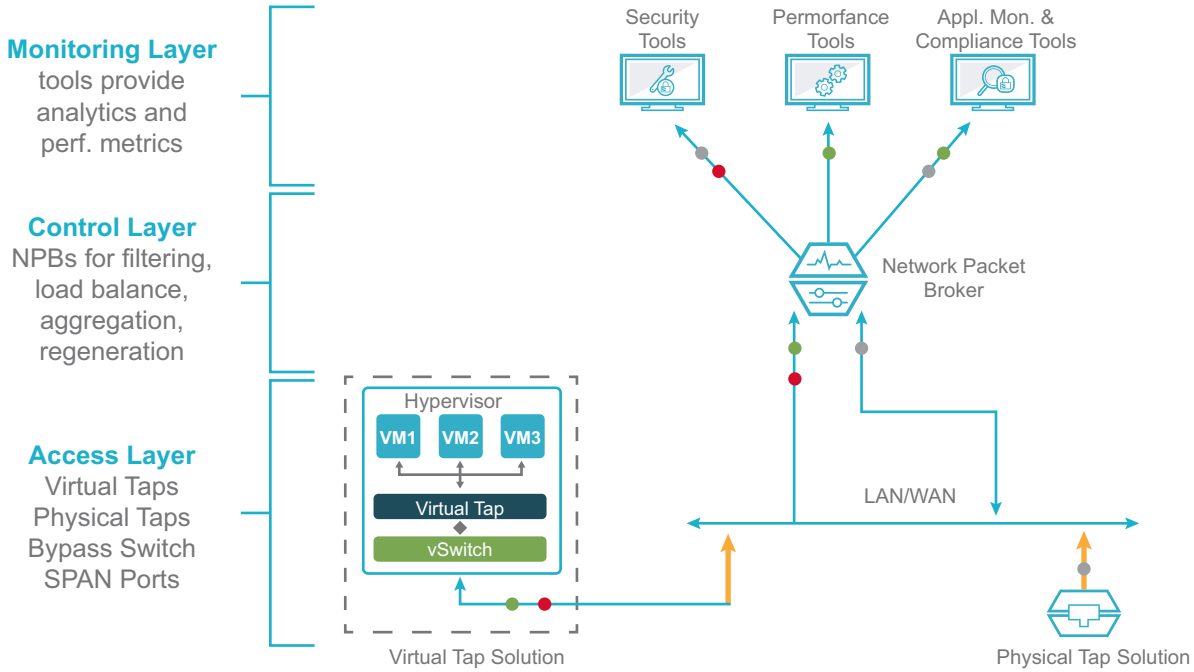

Active vs Standby

PURPOSE OF ACTIVE AND STANDBY

High availability (HA) is achieved through the presence of redundant devices that can back each other up in the event that one of the items fails. The earliest implementations of high availability focused on having a second device in standby mode, sort of like an understudy waiting to step in should the principal dancer be unable to perform. While more than two systems can be involved, the most common implementation is still a primary and secondary node.

The Active/Standby mode of operation (A/S) evolved from initially having the hardware available and connected into the network, to having the necessary software completely installed and available on the second node as well. This is sometimes referred to as a 'warm standby' or 'warm spare.' The next level up is not only having the software installed, but also having the data stored in a way that it is immediately available to the secondary node as well. This is referred to as a 'hot standby.' The cost of a hot standby is a bit more, but the cost of downtime to an organization can be substantial, so most standby configurations today are of the 'hot' variety. Most hot standby implementations recover in some fraction of a minute or seconds.

Modern environments, however, don't tolerate delay very well—even if it's just seconds. Customers, clients, employees, and partners demand better performance. This has increased the interest in having the secondary node move from being a standby to being an active participant in normal operations, to effectively eliminate any delay in recovery. It's like having the understudy on stage mirroring all of the movements of the principal dancer. In IT you see this referred to as Active/Active mode (A/A).

TYPICAL USE CASES AND BENEFITS

Creating resiliency is a part of every IT implementation. With the growth of cyberattacks and the rising cost of data breaches, ensuring the resiliency of network security is more important than ever and a key use of A/A and A/S configurations.

Network Paths

Good network design involves creating redundant paths to maintain operations in the event of a path failure. Alternate paths can be configured to provide visibility to traffic in normal operations (active mode) or to begin operating only if the primary path stops or slows to an unacceptable rate (standby mode). All network paths, whether in A/S or A/A, must have visibility solutions attached, to ensure there are no blind spots where attacks could pass through unnoticed.

Security Monitoring Tools

Many security monitoring tools and appliances are configured in pairs to achieve HA and minimize downtime in the event of a hardware failure. Tool vendors generally require only a single license to operate an HA pair and the second node is deployed as a 'warm standby' that is activated only in the event the primary device goes offline. When that occurs, an alarm is activated to alert security personnel to the failover.

Security Fabric with Network Packet Brokers

Redundancy is also key in a security fabric design, where network packet brokers (NPBs) filter packets and perform preliminary processing of live traffic before passing to the security tools. NPBs keep security tools—many of which are major capital investments—working efficiently and perform critical tasks such as decryption and load balancing. Ixia is the only vendor that offers NPBs with the capability of being configured in A/A mode, using a dedicated HA link for complete synchronization. Both nodes are actively working and each node is aware of the traffic being processed on the other at all times. In the event of an NPB failure, recovery is instantaneous, with near-zero packet loss.

CONSIDERATIONS WHEN EVALUATING FAILOVER OPTIONS

When evaluating your options for failover in your network visibility and security architecture, consider the following:

Speed of Recovery

As mentioned earlier, if the speed of recovery in a failover situation is key, an A/A configuration provides the fastest possible recovery. Both NPBs are actively engaged in traffic processing and there is no time required to transfer processing to the standby node. This can be key to inline security processing, where tools are actively inspecting traffic in real-time. Other applications—such as out-of-band traffic analysis—can accommodate more latency in failover and may be served well enough by an A/S configuration.

Tolerance for Risk

An A/S configuration introduces some additional risk that the standby device has failed silently and will not be able to take over processing in the event the primary device fails. In highly-sensitive environments, this could be an important reason to operate in A/A mode.

Budget Restrictions

The cost of redundant architecture can be a substantial portion of the overall security budget. Many organizations have found it easier to justify the cost of a redundant NPB if it can be put to use immediately in normal operations as an active node. It is the duty of the security team, however, to make sure usage of the two NPBs combined do not exceed 50% of total, because one NPB must be able to completely handle all of the traffic, in order to provide complete failover.

Maintenance without Disruption

In an A/A configuration, maintenance can be completed without any special scheduling and without disruption to the network. Traffic is temporarily routed through the backup devices with no noticeable impact to users. In fact, periodic maintenance proves that the existing configuration is still able to handle the volume of traffic being processed. Any slowness experienced during a maintenance activity would indicate the need to scale the system.

Port Redundancy

The chance of a single port on an NPB failing is much greater than the entire device. Therefore, even in an A/S configuration, you may want to wire your inline tools to the Active NPB on redundant ports, so the secondary port can take over in the event the primary port fails, preventing a complete device failover. This is less important in an A/A configuration since the workload is shared.

Whatever failover mode works best for your environment, make it a priority. Lack of visibility in a complex system puts the entire infrastructure at risk. You want to be able to intervene quickly, accurately, and effectively to protect your network.

Application Intelligence

What do people mean when they talk about application intelligence? For those that haven't heard about application intelligence, this technology (using context-aware data processing) is available through certain network packet brokers (NPBs). It's an extended functionality allows you to go beyond Layer 2 through 4 (of the OSI model) packet filtering to reach all the way into Layer 7 (the application layer) of the packet data. The benefit here is that rich data about the behavior and location of users and applications can be created and exported in any format needed – raw packets, filtered packets, or NetFlow information.

IT teams can identify hidden network applications, mitigate network security threats from rogue applications and user types, and reduce network outages and/or improve network performance due to application data information. Distinct signatures for known and unknown applications can be identified, captured and passed on to specialized monitoring tools to provide network managers a complete view of their network.

TYPICAL USE CASES AND BENEFITS

A powerful use case for application filtering is to improve security tool efficiency. Application filtering effectively allows you to create an early warning system for real-time vigilance. In the context of improving network security, context awareness can provide the following benefits:

- Identify suspicious/unknown applications on the network by exposing the applications that are running on the network. This feature is often an eye opener for IT teams as they are usually surprised to find out that there are actually applications on their network they knew nothing about.
- Identify suspicious behavior by correlating indicators of compromise with geographic location and known bad sites. For instance, maybe there is a user in North Korea that is hitting an FTP server in Dallas, TX and transferring files off network. If you have no authorized users in North Korea, this should be treated as highly suspicious. See this solution brief for an illustration.
- Improve security and monitoring tool efficiencies by flagging trusted data (voice, video, etc.) and allowing it to bypass security tool inspection (e.g. IDS or tool). This can increase the efficiency of security tools by up to 35%. See this case study for an example.
- Improve security by providing immediate, cost-effective SSL decryption activities so that data can be sent clear channel to security tools for inspection and analysis for potential threats.

A second category of benefits include improved troubleshooting and performance capabilities. For this category, application awareness can be used to:

- Find application failures as soon as they happen simply by looking at a dashboard showing all applications running on the network
- Create empirical data to identify bandwidth usage, trending, and growth needs. This allows IT admins to be proactive in managing their resources and forecast expansions.
- Identify new user applications consuming network resources to prevent capacity overload (i.e. explosions) and network outage problems
- Geolocation capability can be used to help quickly locate geographic outages and potentially narrow troubleshooting efforts to specific vendors that may be causing network disruptions. This reduces troubleshooting costs and improves customer Quality of Experience. See this solution brief for an example.
- Find application traffic behaviors that indicate changes in your customers' patterns that you would like to know. As an example, if you are a cable provider who also offers internet service, you would probably like to track the use of your own VoD service vs. competitors like Vudu and NetFlix. This let's you know when a new competitor pops up.

A third category of benefits include improved regulatory compliance capabilities. For this third category, context awareness helps:

- Identify prohibited applications that may be running on your network
- Audit your network policies and usage of those policies. Maybe users are transferring files off network to Drop Box. Maybe employees are using web-based email instead of the official corporate email system; that is linked to anti-viral software for attachment inspection. Without the anti-viral inspection, harmful files can be downloaded onto the corporate network.
- Data masking can be used to hide sensitive data like credit card information before the data is sent to monitoring and logging tools
- Features like Regex can be used to create search strings for sensitive. This allows only relevant data to be sent on monitoring tools for analysis.

CONSIDERATIONS WHEN RESEARCHING APPLICATION INTELLIGENCE AND CONTEXT AWARENESS

When investigating application intelligence solutions, there are several items to consider. Here is a short list of common items:

Extensive functionality

You will want lots of functionality. Once you start using context awareness, you will find more and more needs for it. This technology is an integral part of security, performance, troubleshooting, and compliance initiatives. Look for a solution that is flexible enough to use for current and future needs. A forklift upgrade will be expensive in the future if you find that the current context awareness features won't suit your future needs.

Ease of Use

This will be a critical component. You need an interface that is powerful but intuitively obvious to use. Look for a solution that uses a drag and drop GUI. A command line interface (CLI) will take you 10 times (or more) longer than a drag and drop interface to configure application filters. Look for a vendor that has lots of predefined application signatures. The last thing you want is to have to spend a lot of time creating your own signature definitions.

Technology

A third consideration is around technology. Vendors can say they support context awareness and NetFlow but what features do they really support? Make sure the vendor supports the core features that enable improved security and troubleshooting. Some examples of these features include: geolocation, SSL decryption, data masking, Regex commands, browser and user device type identification, BGP AS, etc.

MORE INFORMATION ON APPLICATION INTELLIGENCE AND CONTEXT AWARENESS

When all components of a visibility architecture are combined, they eliminate the blind spots within your network that are harboring potential application performance and security issues.

Ixia's ATI Processor can deliver context aware information like geo-location, browser type, and device type. The Ixia solution delivers critical intelligence to reduce troubleshooting costs and boost network security protection (especially for indicators of compromise). In addition, Ixia has created several hundred application definitions that can be used to filter application data and forward that data on to appropriate monitoring tools. In addition to the application definition, Ixia can also distinguish application sub-functionality. Using Pandora as an example, the Ixia ATI Processor can understand sub-functions within that application which include "play", "skip", "pause", etc. Information granularity like this reduces application troubleshooting costs and allows you to optimize customer quality of experience.

Blind Spots

What do people mean when they talk about network blind spots? And are these really that important? The answer to the second question is overwhelmingly yes. Blind spots directly correlate to network problems and outages, increased network security risk, and potential regulatory compliance issues. In regards to the first question asking what do we mean by "blind spots", blind spots are hidden reasons for a lack of network visibility.

COMMON BLIND SPOT EXAMPLES

Let's look at some examples of blind spots. Here's a library list (although it's not all inclusive) of examples that I've compiled. While several of these may not apply to your organization, some probably do, right? Scan the list and see if anything matches your network.

- **Silo IT organizations** – Security, Network IT, and Compliance groups don't often talk or share data and can form silos in an enterprise. This can lead to inconsistent data and compliance policies, SPAN port contention issues, improper SPAN port programming that results in incorrect or missing data captures, and plain old data conflicts that arise from collecting data at the wrong places.

- **Use of virtualization technology** – According to Gartner Research, up to 80% of virtualized data center traffic is east-west (i.e. inter- and intra-virtual machine traffic) so it never reaches the top of the rack where it can be monitored by traditional tap and SPAN technology. While virtual tap technology exists to counter this threat, according to the Ixia 2015 virtualization study, 51% of IT personnel don't know about the technology. For instance, an Ixia healthcare insurance provider had zero visibility into 100+ virtual hosts. This was immediately solved when they installed the Ixia Phantom vTap.

- **SPAN port overloading** – An Ixia case study shows the various problems that a national pharmacy ran into with SPAN port contention problems and the fact their SPAN ports were also dropping packets and creating a loss of information due to a data overload condition. SPAN ports can, and will, drop monitoring data if the CPU is overloaded. Besides running into port contention problems, the case study also shows that the customer ran into problems splitting and filtering the data from the SPAN ports. For more information on general SPAN port visibility issues, see this article.

- **Rogue IT** – When users add their own Ethernet switches, access points (from an iPhone), use offsite data storage (like Box), or add something else to the network, company security policies are often subverted which opens the door to security, compliance and liability issues. IT rarely knows anything about these devices, especially as they can appear sporadically, like Wi-Fi hot spots.

- **Mergers and acquisitions** – The blending of disparate equipment and systems often causes interoperability issues which adds to system/application downtime, system capabilities being turned off to improve network performance, and the scaling back/elimination of network and application monitoring while extensive network re-architecting takes place. This results in very limited visibility (i.e. blind spots) because no one really knows what is happening.
- **Addition of new network equipment** – When new equipment is added, there is often no record as to who owns it and what it does. The equipment can get "lost" and forgotten about, especially if IT key personnel leave the company or change departments. "Lost" equipment that is still functioning in the network can be a source of security vulnerabilities due to lack of proper software updates and unknown user access privileges.
- **New equipment complexity** – New equipment is often complex to understand, i.e. what it does and how best to use it. For data networks, complexity never seems to take a rest at all. The rate of increase of this complexity has been characterized by David Cappuccio at Gartner who stated in a Gartner Symposium back in late 2012 that for every 25% increase in functionality of a system, there is a 100% increase in complexity. See the blog that Eric Savitz (with Forbes) wrote about that symposium. If IT doesn't have time to do the research on new equipment and how to properly program it, they often stop using the equipment and then eventually forget about it. The equipment can often remain running in the network even though it isn't being utilized.
- **Network complexity** – When new links and office locations are added, they can be set up with different VLANs, sub nets, etc. to geographically segment them. These segmented networks often have separate equipment that is used for remote logon, authentication, etc. that makes it hard to track what is happening at those locations.
- **Inconsistent monitoring/data collection policies** – This can occur from multiple sources but one of the common effects is that virtual monitoring equipment policies and physical equipment monitoring policies are often different, which can cause compliance data mismatch, requisite data that is simply not captured, and security issues. See this case study for an example.
- **Network planning issues** – In many cases, the requisite data just doesn't exist at all. This can be a common experiences for organizations with external customers. For instance, service providers (especially wireless service providers) need good customer data (service holes, malfunctioning radios, poor coverage, and even customer dissatisfaction) to properly plan their networks and deliver a better quality of experience.
- **Network upgrades that are postponed** – Postponing upgrades can result in continuing to use old and antiquated equipment that has limited uses on a higher speed network. Network performance then becomes slow, which affects IT's ability to solve problems as fast as required. More information is available in this ebook on upgrade issues.

- **Network upgrades are implemented** – Just the action of necessary upgrades can result in blind spots. One example is if new higher speed equipment is added. This equipment may end up overloading various components of the network, especially monitoring and security tools, with too much data. This is especially true if any monitoring and performance tools weren't upgraded at the same time. These tools can become overloaded and lose (i.e. drop) data or overwrite buffers/logs at a faster than expected pace. In addition, tool dashboards are often limited in what they can see which allows the blind spot to remain hidden. More information is available in this ebook on upgrade issues and common vulnerabilities can be found in the CVE database.

- **Addition of new applications** – A common blind spot for hospitals is access to application data and application performance trending. In this case study, the customer was using the EpicCare Ambulatory Electronic Medical Record (EMR) application from Epic but was having problems correlating all of the information from their different systems.

- **Security and network audits are postponed or rarely occur** – This action will often result in a safe and cozy harbor for various threats and malware on your network. It's hard to say what will be hidden but whatever it is, I'm sure you don't want it. See this resource for more information.

- **Anomalies** – Unexplained network events happen and are often addressed by IT but if they are spurious and random in nature, and they go undiagnosed, this can result in larger problems later on. Ixia has several customers who have eliminated their network anomalies and also realized a mean time to repair (MTTR) reduction by up to 80%.

- **Incorrect equipment programming rules** – An example of this is firewall programming, which is rules-based and typically processed through access lists. When the traffic matches a rule, it is immediately forwarded on, even if subsequent rules exist to tailor the information. This can cause gaps in network security because the packet was routed before the correct security tool got to see the correct information.

HOW TO ELIMINATE BLIND SPOTS

So, when it comes to your specific network, where are your potential blind spots? If some of the blind spots listed above apply to you, you've typically got two ways to respond – either in a proactive or reactive manner. The reactive approach is straight forward, just wait until something happens and then go fix it. While it's the simplest approach, it's also usually the costliest in terms of locating exactly what issue the blind spot caused (which usually increases your mean time to repair). In addition, it often necessitates the purchase/implementation of expensive long term fixes or multiple "Band-Aid" fixes that never really "fix" the problem. In any case, this approach is very straight forward.

If you want to follow a proactive approach, the best solution is to design a visibility architecture. This involves more upfront cost and planning but will normally pay for itself very quickly. The visibility architecture is a plan you create for organizing exactly how you want your monitoring tools to connect to the network. This involves how they connect (taps or SPAN ports), where they connect (edge, core, which branches, etc.), and how you groom the monitoring data before you send the stream to a tool (packet filtering, application filtering, deduplication, packet trimming, decryption, aggregation, etc.). If you want to learn more about designing a visibility architecture, check out this whitepaper.

To end blind spots in your network, you need to be able to see everything. Unknown issues and "soon to be problems" exist in every network to some degree. To achieve the goal of ending blind spots in your network, you'll need to implement a visibility architecture. It's not hard or complicated, but it does require some planning. At the same time, the sooner you can accomplish this step, the faster you can integrate a visibility architecture with your IT network. And the sooner you can realize cost and productivity savings.

CONSIDERATIONS WHEN RESEARCHING VISIBILITY ARCHITECTURES

When considering visibility architectures, there are several items to investigate. Here is a short list of common items:

Flexibility, i.e. choice

You will want, and need, options. This includes the flexibility to deploy inline and out-of-band visibility solutions. It also includes the ability to monitor your physical and virtual data center traffic. Application Intelligence is another area to look for. While you may not want to engage in all of these activities right away, you should look for a solution that allows you to add the pieces you want, when you want, without a forklift upgrade.

Ease of Use

This will be a critical component that will heavily influence your total cost of ownership (TCO). Look for a solution that uses a drag and drop GUI interface. A command line interface will take you 10 times (or more) longer than a drag and drop interface to configure filters. The management system should also be able to handle everything—from global element management, to policy and configuration management, to data center automation and orchestration management. Engineering flexible management for network components will be a determining factor in how well your network scales.

Technology

A third consideration is around the technology. Buyer beware applies to this market place (just like others you are used to). While vendor products may sound the same, they usually aren't. In general, a strong consideration should be to purchase NPBs that run at line rate under all conditions. Only a very few NPBs do this. Anything less adds delay to your monitoring effort. For inline solutions, this line rate will be absolutely critical. You will also want failover technology that is as fast as possible for inline solutions.

Data Access

Data access is another area of concern. Consider using Taps instead of SPAN ports for your data access technology. Taps are superior to SPANs for several reasons, see this analysis. One key difference is that SPANs provide summarized data (instead of a complete copy of all data) that can often be missing key data you need for proper problem resolution. Another area to investigate is whether your tools need packet data or NetFlow data. One last thing to consider is if your tools need additional data from application intelligence functions to further improve their performance.

Bypass Switches

Ups and downs are part of life. But when it comes to your network, frequent, or long periods of downtime simply aren't acceptable.

The latest inline security tools can prevent network downtime by keeping threats at bay. Yet they come with risks too. Inline tools are single points of failure (SPOF) in your network, meaning that if they fail, the whole network fails with them. When inline tools stop working, or lose power, or get congested, network traffic can no longer flow through them. So your network, and your business grinds to a halt. Costing you precious time and money.

But bypass switches can save you from this expensive eventuality.

WHAT IS A BYPASS SWITCH?

A bypass switch is a simple piece of hardware that allows you to connect inline security tools to your network – without the risk of network downtime.

Picture the way a car accident causes road congestion and stops the flow of vehicles. Typically, the police will step in to redirect oncoming traffic, so as to avoid a huge pile up, and keep the stream of cars moving. Bypass switches prevent outages and keep network traffic moving in a similar way.

When inline tools fail, bypass switches automatically kick in. They redirect network traffic so that it flows around the failed tool, instead of through it. The network traffic is able to "pass by" the blockage caused by the tool. So your network stays up and running – even if the tool isn't.

Bypass switches can detect when an inline tool has failed or lost power through heartbeat packets. Heartbeat packets are signals sent from the bypass switch, through the inline tool at regular intervals. If a packet doesn't make it back to the bypass switch, the inline tool is assumed to have failed, and network traffic is rerouted.

INTERNAL VS. EXTERNAL BYPASS SWITCHES

Some inline security tools come with ready-to-go internal bypass switches. On the surface, this sounds ideal. Why pay for a standalone device, when you can get inline tools pre-fitted with bypass switch technology? External bypass switches have a number of advantages over internal bypass switches.

CONSIDERATIONS WHEN RESEARCHING BYPASS SWITCHES

Here is a short list of items to research on bypass switches:

Simplicity and Network Reliability

At some point, your inline security tool will need to be taken offline, and physically removed from the network. This may be for planned maintenance or upgrade purposes. But it could also be due to an unplanned event. If the tool has been installed with an external bypass switch, it can be removed without impacting the flow of network traffic.

The image below shows how an Intrusion Prevention System (IPS) – a type of inline security tool – can be disconnected without breaking any network links.

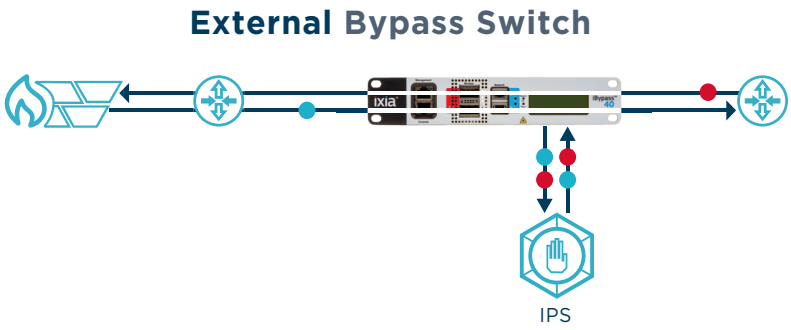

In contrast, removing an IPS with an internal bypass switch means that network links must be broken:

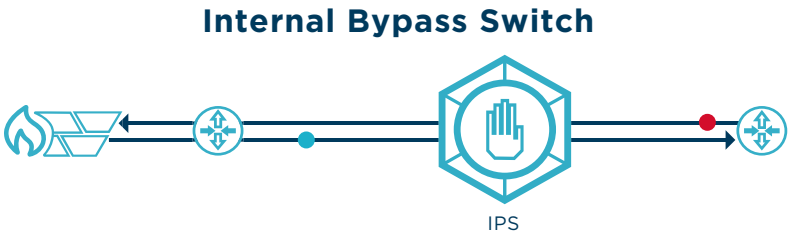

This creates costly downtime and service interruption for your organization.

Security

Since external bypass switches are purpose-built, they're better equipped to keep your network safe. Internal bypass switches come as an "added extra" to inline tools. So they may not have all the features of an external solution. For example, external bypass switches can redirect network traffic around a failed tool, while also notifying network / security teams of the blockage so swift action is taken.

Purpose

An inline tool that includes integrated bypass capabilities (i.e. a jack-of-all-trades) simply won't be as effective as specialist equipment.

Command Line Interface

Computers do amazing things. And it doesn't take a tech expert to know we've only scratched the surface. With driverless cars, 5G technology, and job automation on the horizon, computers will soon be doing more to support and sustain our way of life than ever before. But, on a very basic level, this all starts with one thing—we must be able to tell computers what we want them to do.

WHAT IS A COMMAND LINE INTERFACE (CLI)?

Let's first define what a user interface (UI) is. A user interface is a system by which a user and a computer interact—swapping information and instructions. A CLI is a type of user interface. Other types of UI include menu-driven interfaces, and graphical user interfaces (GUI).

Command line interfaces work by allowing the user to issue commands to the computer in the form of lines of text. Users must be familiar with a vocabulary of specific text commands that the computer will recognize and respond to. The user keys in the relevant command, and presses the 'enter' / 'return' button. The computer processes the command, and the result is displayed on screen. MS-DOS is a well-known example of a CLI system.

TYPICAL USE CASES

CLI began life as the primary way in which people interacted with computer systems. Today, the graphical user interface (GUI) has surpassed CLI in popularity, but CLI continues to be used. This is because there can be benefits to using CLI. These include:

- CLI requires fewer computing resources than GUI. This means CLI is a good option when there is limited processing power available.
- It's ideal for experts who are familiar with the necessary command vocabulary and who wish to execute (sometimes multiple) tasks at speed
- It's easier to automate processes using CLI

CONSIDERATIONS WHEN USING CLI

There are a number of drawbacks to CLI, which is why a GUI is more commonly used. Some CLI drawbacks include:

GUI interfaces can be a lot faster

If more than a simple command or two is required to implement a function, a GUI is typically faster. Observations from Ixia systems show that the GUI can be four times or more faster.

CLI is not intuitive

CLI is reliant on the user being fluent in command vocabulary. The commands (which are not obvious) must be typed into the computer interface. A GUI removes this burden.

CLI is more error prone than a GUI interface

The cryptic command are very syntax-oriented which means that if anything is out of place when typing the commands, it may result in an error. Debugging errors is difficult and tedious.

CLI is not beginner-friendly

A proper GUI abstracts away the complexity for users which reduces training time and ultimately reduces the total cost of ownership of a monitoring equipment purchase.

CLI is visually sparse

This makes it uninteresting to look at and increases the probability of user fatigue and error.

Networks are already complex

Network Packet Brokers (NPBs), which capture data from across the network and pass it on to the desired analytics, monitoring and security tools, have become the cornerstone of network security. NPBs are powerful, but they require careful management and configuration. NPBs that must be configured using a command line interface, require in-depth training and expertise. They offer little to no ease-of-use. And when IT teams are faced with hard-to-use configuration tools, valuable time, money, and effort can be wasted simply completing basic security tasks. Blind spots and network vulnerabilities become a risk.

MORE INFORMATION

GUI-based NPB configuration tools are more beneficial than their CLI-based counterparts. User-friendly NPBs built around a GUI, require less training, and can be operated by IT administrators of varying levels. They free up IT to take full advantage of NPB capabilities.

An NPB with a proper GUI should deliver the following benefits:

- Reduced capacity for errors, as the drag and drop interface means no risk of mistyping commands
- Reduced need for special training – no new command vocabulary needs to be learned, allowing junior engineers to use Ixia's configuration tools
- Easier, and lower cost filter testing
- Reduced need for troubleshooting

See for yourself how a graphical user interface reduces NPB management complexity by watching this video. Or read more on the tangible benefits of Ixia's user interface over Gigamon's here.

Crash Carts

Have you ever heard the term “crash cart”? One common definition is used by the healthcare industry for a movable cart located in a hospital which contains basic equipment that can be used to help resuscitate patients that have stopped breathing. In the IT world, there is a similar definition. IT crash carts can help troubleshoot and resuscitate networks.

PURPOSE OF CRASH CARTS

An IT crash cart refers to a movable cart that typically has a computer (probably a laptop), one or more protocol analyzers, cables, screw drivers and other tools, and commonly used special purpose tools, etc. This pre-assembled kit typically saves an engineer time, since they don’t have to hunt down equipment when they need it during an emergency. The basic equipment is ready to go whenever they need it.

TYPICAL USE CASES

Crash carts are typically used for troubleshooting network and equipment problems. By placing the essentials on a rolling cart, IT was hoping to improve response times to network issues and outages so they could maximize network uptime. The cart is wheeled to wherever it is needed and then connected into the network or device (tap, switch SPAN port, monitoring tool, computer, or other device) to start collecting data which can then be sent to the laptop and protocol analyzer for analysis and debugging. If the crash cart doesn’t have the right equipment, data is collected and then fed into purpose-built monitoring tools for additional analysis.

CONSIDERATIONS

If you are using crash carts, something you may want to consider is that a visibility architecture using taps and network packet brokers (NPB) could be a better solution for you. While the crash cart process was an improvement over previous processes, crash carts are an outdated process now.

Here are some reasons why:

Faster Troubleshooting Times

Once you have installed taps and an NPB, you will no longer need to hunt for troubleshooting carts or wait on maintenance windows. The troubleshooting equipment should already be connected to the NPB. And by connecting the NPB to multiple taps throughout your network, you have instant access to the data you need. You don’t need to get Change Board approvals to connect into the network (you’re already connected), so you can start debugging problems right away. The alternative is to wait for a Change Board decision that approves the activity but “only at 2 am two weeks from this Tuesday” or something like that which means the problem will exist for way too long.

Improved Reliability

Once you have installed taps and an NPB, you will probably no longer need to touch the network, except as part of a fix for the problem. Whenever the network is touched, there is an inherent danger that the network can be impacted or even brought down through some unintentional complication. This is why most companies implement a Change Board approval process—so that the organization is aware of any potential service interruptions.

Ease of Use

This is a critical component that will heavily influence your total cost of ownership (TCO) and improve troubleshooting response. With the NPB, taps and tools connected together, you can start the debug process immediately, wherever you are. The NPB management system can provide remote access to the NPB, making the solution very easy to use. You can also create packet captures (PCAPs) with an NPB as well.

Data Access

Data access is another area of improvement. With a crash cart, you need to physically move it to wherever it needs to go so you can connect into the network. With the NPB already connected, you don't need to physically relocate anything, unless you don't have the data access you need. In that case, it might be easier for you to just install a tap and connect that to the NPB instead of the crash cart. What about troubleshooting across geographic areas? Crash carts might mean that you have to roll a truck and physically drive to a location to access the network, especially if you have a specific problem and need a specialized tool. This use case is becoming much more common as IT switches from just trying to maximize network uptime to maximizing application availability. A reliable network is now considered table stakes. IT has been focused to make sure that applications are available 24 x 7 x 365 as well. This makes the use of crash carts much harder as application performance testing often requires specialized tools and hours/days worth of troubleshooting data to analyze.

Data Masking

When it comes to network security, full visibility isn't always ideal. There are some things that should stay hidden. Most companies have data compliance requirements they must adhere to. HIPAA, PCI, and internal best-practice policies mean Personally Identifiable Information (PII) must be handled with care. Data masking helps organizations control who has access to this sensitive data.

WHAT IS DATA MASKING?

Data masking is different from restricting data access. Access restriction renders data invisible. Data masking replaces vulnerable, or sensitive, data with information that looks real. When data is masked, it's altered so that the basic format remains the same, but the key values are changed. For example, a long card number could be masked in the following manner:

1234 5678 1234 5678 → 1234 XXXX XXXX 5678.

Data can be masked, or changed in a variety of ways:

- **Substitution** – When external, unrelated, or randomly generated data is used to replace parts of the real data. E.g. a random list of names could be used to replace a real list of customer identities.
- **Shuffling** – When data is swapped or shuffled within the database.
- **Encryption** – When sensitive data is converted into code, using an algorithm. The data can only be decrypted with a key.
- **Number / Date Variance** – Where each number / date in the set is altered by a random percentage of its real value.
- **Masking Out** – Where certain fields or parts of the data are replaced with a mask character (an X, for example).

Whichever method is used, the data must be altered in such a way that the original values cannot be obtained through reverse engineering.

TYPICAL USE CASES

On the whole, data masking helps organizations comply with data protection requirements. Here are some specific situations in which data masking capabilities are particularly useful.

- **Testing / Outsourced Data** – Companies may need to provide true-to-life datasets to develop relevant software. In less secure test environments such as this, realistic data is needed, but companies need not risk data security by using actual customer information. Since creating a 'fake', but realistic dataset from scratch is time-consuming, companies can use the real data they have, but alter it through data masking to protect customers.

Similarly, if business IT operations are outsourced to another organization, data can be masked to prevent exposure of real data to more people than necessary.

- **Network Data** – Organizations often need to record and monitor network data. But compliance requirements mean they must avoid storing PII. With data masking, companies can record network data while hiding sensitive data.
- **SSL Decryption** – Secure Socket Layer (SSL) encryption is the standard technology used to send private information. For security purposes, organizations must decrypt and examine any SSL traffic on their networks. One of the dangers of SSL decryption is that it makes sensitive data available to anyone with access to network monitoring tools. Clever network tools can decrypt SSL data, while masking data that doesn't need to be exposed.

CONSIDERATIONS

According to research survey conducted by Enterprise Management Associates in 2016, data masking is one of the Top 5 most commonly used packet broker features.
So, if you are considering a purchase for data masking equipment, here are some things to keep in mind:

Use of masked data

Make sure you understand how you intend to use the data before you make your purchase, as this can save you time and money. For instance, is the plan to simply distribute the data to a data loss prevention (DLP) device for analysis or do you plan to access the data natively for searches? If you plan to access the data natively, then your solution needs to support regular expression (Regex) search capability. Once the specific information, or type of information, is found that matches the search criteria, that data can be sent to a tool (like a DLP) for further processing. This search capability allows your monitoring tools to be more effective, as they have less data to sift through.

Easy access to the data

If you need access to the data for Regex searches, consider purchasing a network packet broker that supports data masking. This will allow you to easily collect the data, search through it, and then forward it to monitoring equipment (like a DLP) for advanced data searches, data storage equipment (like a SAN), or compliance and recording tools.

Distribution of masked data

Once the data is masked, how do you plan to distribute it to the appropriate monitoring tools? This is where a packet broker will come in handy for you to distribute the data to the device(s) that it needs to go to.

MORE INFORMATION ON DATA MASKING

Data masking is a powerful capability that can help your regulatory compliance initiatives. When recording network data, it allows you hide sensitive information with ease.

Read more about Ixia's ATI Processor data masking capabilities and network packet broker (NPB) solutions to see how simple and easy data masking can be.

Fail Closed

Simply stated, failing closed is when a device or system is set, either physically or via software, to shut down and prevent further operation when failure conditions are detected. This strategy is common in situations where security concerns override the need for access. We encounter this every day when we forget the password to a seldom-used personal account and are denied entry. A physical example is the failure of a metal detector at the entrance to a federal courthouse, which leads to a long line of people waiting to get in at a second door, while a technician tries to repair the first door. In these situations, access is a second priority to security.

USE CASE & BENEFITS OF FAIL CLOSED

The primary use case for Failing closed is to prioritize network security. In an IP network, security appliances like firewalls can be configured to fail closed which prevents incoming Internet traffic from being passed into your internal network when the firewall is unable to confirm that the packet is allowed. The network outage that results from a firewall outage can be minimal, if a backup firewall quickly takes over processing duties (like the second door at the courthouse). The fail closed condition generally provides greater confidence that a cyber threat or attack will not sneak in while a firewall is offline.

CONSIDERATIONS

It's important to note that the fail closed strategy, even for a device like a firewall, has not always been the rule. In some environments, network interruption can be a greater concern than security, leading to the choice to fail open. This was more frequently the case in the early days of firewall deployment, when organizations were learning how to balance the need for security inspection with network availability.

Fail Open

A system set to fail open does not shut down when failure conditions are present. Instead, the system remains "open" and operations continue as if the system were not even in place.

This strategy is used when access is deemed more important that authentication. Healthcare systems are sometimes operated on a fail open basis, such as when emergency care is provided even without authentication of insurance coverage or the ability to pay. The risk (of non-payment in this case) is essentially mitigated by performing authentication after-the-fact.

Another example often cited is when a door with an electronic locking mechanism is automatically unlocked when the system fails and is unable to authenticate access credentials. This ensures an exit is made available, particularly in the event of a fire or natural disaster that disables electronic systems.

USE CASE & BENEFITS OF FAIL OPEN

Protect network availability, i.e. network uptime

Historically, some organizations considered inline deployment of a network firewall to be a "nice-to-have," rather than an essential element of IT security. When a firewall failed, they preferred to have it fail open and let Internet traffic proceed on into the internal network without authentication. The thinking was that, the majority of traffic was safe and the risk of a network breach was low, so it did not make good business sense to interrupt network operations. The business risk was minimized by prioritizing firewall restoration to limit potential exposure and by analyzing copies of network traffic (using out-of-band tools) to detect suspicious activity after the fact. The fail open condition prevailed in situations where access was deemed more important than security.

Supplement another security appliances

There are security solutions that organizations may want to operate in a fail open condition to supplement the function of existing security appliances. One example is an advanced malware protection (AMP) sandbox, which is used to execute unknown files in a safe environment and provide the results to anti-malware solutions. Since the sandbox is supplementing the main device, it's failure may not require a complete shutdown of processing.

Deployment and testing

Another practical use for fail open is during the initial deployment and testing period of a new security appliance. Configuring a new device to fail open allows the team to become comfortable with the operation and learn how to respond to alert situations without becoming overwhelmed. Once the team feels confident, the device can be switched over to a fail closed condition, for greater risk management.

CONSIDERATIONS

The primary consideration for a fail open situation really comes down to what is the design of your security architecture? This architecture will have a lot of impact on your choice.

Fail Safe

Fail safe refers to a device that is configured to protect all other components in the system from failure, in the event the device itself fails. Practically, this can have the same result as failing open, but fail safe is often achieved through addition of a separate device, known as a bypass switch.

TYPICAL USE CASE & BENEFITS

Bypass switches are deployed "in front of" network devices and work by establishing a direct connection to the device and monitoring its ability to receive and process traffic. This is achieved by sending a very small network packet, called a heartbeat packet, to the device at very fast intervals—generally one every couple microseconds. If the packet is returned, the bypass remains open; if the packet is not returned, traffic is bypassed around the device and moved along to the next switch in the network.

CONSIDERATIONS

Many network security appliances, such as next generation firewalls and IPS solutions, now include an internal bypass function. However, internal bypasses do not provide all of the functionality of an external bypass switch.

An external bypass switch deployed in front of a network device can be activated proactively by the IT staff, to take a device offline for regular maintenance, periodic troubleshooting, or repositioning in the network. The external bypass essentially removes a particular device temporarily from the active network, eliminating the need to wait for a network maintenance window to perform upgrades or respond to support issues.

Failover

Failover is the ability to recover the functionality of network devices that fail. This is a broader concept than fail safe, which only specifies only no adverse impact to other components. Failover implies recovery of functionality, achieved through redundancy.

TYPICAL USE CASE & BENEFITS

This capability is often integrated in to core pieces of equipment, e.g. network switches and packet brokers, to ensure maximum network uptime. While the capability can be including into security and monitoring tools themselves, many failure scenarios make the practicality of that solution null and void. This functionality is better off being part of a purpose-built solution.

CONSIDERATIONS

External bypass switches and network packet brokers are now available with the ability to designate an alternative path for traffic in the event of a network device failure. For example, should the primary IPS appliance fail, when the external bypass switch or network packet broker detects the failure (within microseconds of the event), the switch can automatically begin sending traffic to a secondary, backup appliance. This can be a cost-effective solution for achieving resiliency.

Floating Filters

An interesting troubleshooting benefit of certain network packet brokers (NPBs) is the ability to create unassigned data filters. At Ixia, we call these floating filters because they aren't attached to any network port, so they are free floating. The advantage of this type of filter is that it is already pre-configured in the system. So when needed, the time it takes to actually deploy the filter is almost negligible.

PURPOSE OF FLOATING FILTERS

Floating filters are only partially connected within the network packet broker (NPB) programming. They are typically connected to a specific tool but not to a network port. The power of the floating filter is that it is already created and connected on the tool side. When needed, the tools can instantly be connected to a network port to analyze incoming data. This speeds up diagnosis time since the forensic tools are already in place and in standby mode.

When you're in a troubleshooting situation, minutes matter. According to the 2016 Cost of Data Center Outages study conducted by the Ponemon Institute, the average cost of a data center outage is $740,357 and lasts for about 95 minutes. This results in a cost of $7,793 per minute of downtime. A rapid response is needed to control costs. Since the floating filter is already created, this can save you several minutes, especially when compared to configuring filters manually using CLI.

To activate the filter takes less than 1 minute. You just draw a connection from the network port to the floating filter. It's that easy. If you need to make any filter adjustments, they are simple button clicks.

TYPICAL USE CASES

A typical use case would involve using something like a Wireshark tool or some sort of protocol analyzer. You can also use this technique for tools that you use often. For instance, maybe you've got a commonly reoccurring problem happening but you don't have a long term fix for it yet. Any tool that is used often can be set up with a floating filter and pre-staged for problems.

Another use case is to have a senior engineer create various types of filters and save them to the filter library. This increases speed of deployment and the accuracy of the various data filters created. In addition, the floating filters can be connected remotely by using the packet broker management system when needed. This gives you 24 x 7 x 365 diagnosis capabilities from remote locations.

CONSIDERATIONS

Here are some things to keep in mind regarding floating filters because most vendors do not support this valuable capability:

Faster Troubleshooting Times

here is no “mapping” or extensive configurations needed. The power of this feature, is to pre-stage specific troubleshooting filters and connect them to standby troubleshooting tools like protocol analyzers (like Wireshark, etc.). This is what allows you to dramatically cut data collection times because it literally only takes a minute or less to connect the pre-existing filter to a network port.

Better Accuracy

These types of filters are predefined and can be stored in a library. This means that someone on the team who is an expert for certain activities can create special purpose filters that the whole IT team can use. Since the filter is created and validated by the company expert, there will be less issues with filter accuracy and correctness should a junior engineer, or someone not familiar with all of the nuances with a particular type of test procedure, need to conduct specific types of network testing.

Role-based Permissions

Filters should have the ability to be locked down, i.e. allow role-based permissions. This allows an individual or group to be able to access certain filters but not everyone. You may want filters that can be accessed by everyone but at other times, you want the ability to lock a filter so you do not have to check it all the time to make sure no one changed the core filter programming. When an event happens, you want to start data captures, packet captures (PCAPs), and data analysis as fast possible. By locking a filter down, you have the peace of mind it has not been changed and is “good to go.”

Ease of Use

If any configuration of the filter is needed for some reason, it should be as simple as making mouse clicks. A good packet broker will display the existing filter within the main User Interface so it’s easy to see the connections and easy to understand what a particular filter is used for. You can drag and drop to start the flow of data to the filter. The packet broker should support a remote interface so that you can make changes or place the floating filter into service remotely, i.e. no need to drive into the office.

MORE INFORMATION

When all the components of a visibility architecture are combined, they eliminate the blind spots within your network and make troubleshooting much easier and faster. Floating filters allow you to improve your trouble alert response times.

Inline and Out-of-Band

When it comes to network monitoring, there are two scenarios—out-of-band and inline (in band). This definition typically refers to the placement of the equipment from the monitoring tool's perspective. Basically, is the monitoring tool in the critical path of network data or not? If the tool is not in the main data path and just using copies of the packets, then it is called out-of-band. If it is actually processing the original data, it is said to be inline. It's that easy.

The next question, of course, is why does it matter?

PURPOSE OF INLINE AND OUT-OF-BAND MONITORING

The type of monitoring scenario, out-of-band or inline, effects the placement of monitoring equipment, the type of equipment used, and the monitoring activities you can conduct as part of your visibility architecture. For instance, firewalls are typically located at the company's main network interface to the outside world. Because of this, they are placed inline. An intrusion detection system (IDS) is typically not placed inline. It is installed as part of an out-of-band scenario since, while it is used to sample data for intrusions, it is not intended to inspect every packet that traverses the network.

For inline tools, data access starts with a bypass switch. Think of this as a special tap for monitoring tools that you insert directly into the flow of the network data. If you had just inserted the tool, and it failed completely or you yanked the tool out, this would directly affect, i.e. stop the flow of data to the rest of the network. A bypass switch has fail-over capability that let's the network survive if the tool connected to it fails. If a network packet broker (NPB) is inserted between the bypass switch and the tool(s), then other functions like filtering and load balancing of network data are possible.

Here is a simple example of an inline scenario:

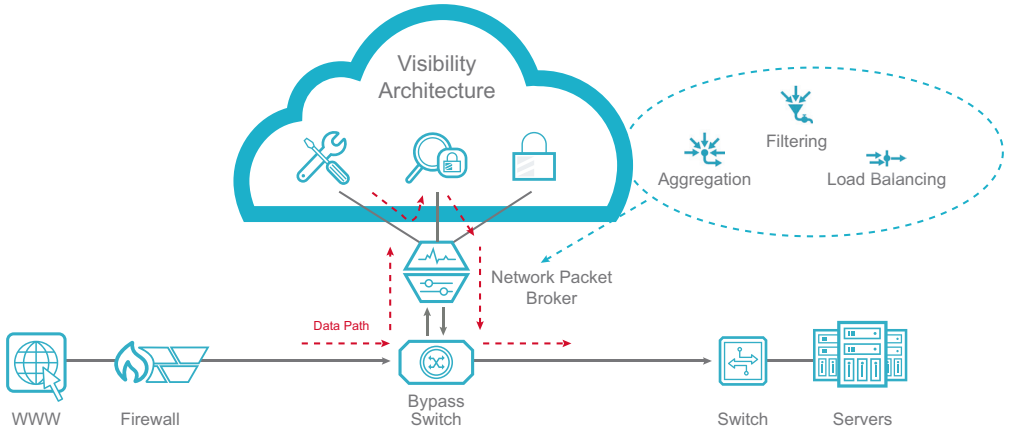

In an out-of-band monitoring scenario, a passive tap is inserted into the network for data access. This device doesn't need the fail-over capability because the monitoring equipment is not directly in the flow of network traffic, so it's simpler. In fact, it's essentially set and forget. You normally don't have to do any programming to taps. While the tap is in the direct path of traffic, all equipment that is receiving a copy of the traffic is completely out of the traffic path for network traffic. You can connect or disconnect whatever equipment you want to the tap monitoring port and it will NOT affect the rest of the network. In this scenario, a packet broker can also be inserted between the tap and tool to perform filtering, load balancing, deduplication, packet slicing, data masking, and lots of other functions.

Here is a simple example of an out-of-band scenario:

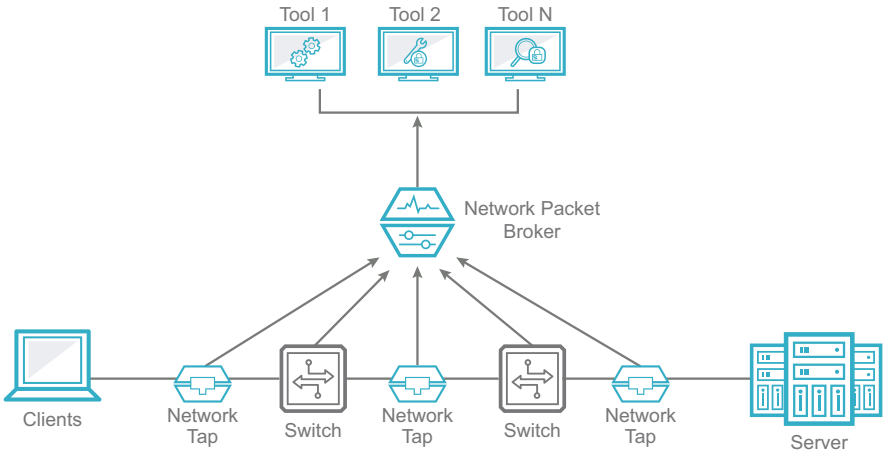

TYPICAL USE CASES

Here are some common use cases for inline and out-of-band monitoring scenarios:

- **Security (both scenarios)** – Security monitoring solutions involve inline components like firewalls, intrusion prevention systems (IPS), honey pots, serial chaining of suspect data, and threat detection solutions. Out-of-band solution examples include forensic analysis with data loss prevention (DLP) tools, intrusion detection system (IDS) analysis, and forensic packet recording.
- **Cost controls (both scenarios)** – Both scenarios offer cost saving capabilities like load balancing, data filtering/discrimination, floating filter creation, remote management, etc.
- **Survivability (both scenarios)** – Both solutions provide improved survivability like the bypass switch and high availability for inline security tools as well as redundant components and fail-over NPB functionality for out-of-band solutions.
- **Performance monitoring (both scenarios, more common for out-of-band)** – While some performance monitoring tools can be implemented as part of inline scenarios, most of these solutions will be out-of-band and focus on application and network monitoring. Proactive monitoring, the ability to test the network in real-time, is also an out-of-band solution.
- **Application intelligence (both, more common for out-of-band)** – This feature is more common for out-of-band scenarios because of the analysis utility for application data.

Application data can be used to help identify indicators of compromise, proactive troubleshooting, and to improve/demonstrate regulatory compliance.

- **Troubleshooting (out-of-band)** – The out-of-band scenario allows for the collection of various data points that can be used to pinpoint problems. The existence of the data typically doesn't reveal the problem itself. That data needs to be sent to an analysis tool that requires a certain amount of time to analyze the data before a useful conclusion can be made. This time delay necessitates the need for an out-of-band scenario.
- **Compliance (out-of-band)** – The out-of-band scenario allows for data masking and packet slicing to conceal packet data while it is stored. Data can also be filtered and sent to special purpose tools, like logging tools, for data storage to demonstrate compliance to various regulatory standards.
- **Virtual data center monitoring (out-of-band)** – Out-of-band solutions are used for gaining access to monitoring data within a virtual data center. This includes a special purpose tap, called a virtual tap, to capture the requisite data and send that off to monitoring tools for data analysis.

CONSIDERATIONS

Here are some things to keep in mind to help determine whether you need an inline or out-of-band monitoring solution.

Monitoring Goals

What information do you wish to collect from the network and where are you planning to get it from? Determining an inline scenario is usually fairly simple. For instance, do you need to collect and process the actual data packets, or just copies of the data packets? Do you need to analyze and inspect every packet? These are inline scenarios. Out-of-band will essentially be everything else. This includes performance monitoring, forensic analysis of security risks, data analysis for compliance, troubleshooting network issues, etc.

Performance

Performance of the equipment will be paramount. You need taps, bypass switches and packet brokers that process data at full line rate under full load. Some of the visibility solution providers out there sell products that cannot operate at full line rate. So pretest your solution before you buy.

Scalability

Scalability is important for long-term cost controls. The solution needs to be able to support current bandwidth requirements and future requirements. You also need solutions that can be upgraded to higher data rates, like 100 GE, in the future.

Ease of Use

Packet broker filter creation needs to be as simple as making mouse clicks. A good packet broker will display filters within the main User Interface so that it's easy to see the connections and easy to understand what a particular filter is used for. You can use drag and drop functionality to start the flow of data to the filter. The packet broker should also support a remote interface so that you can make changes or place the floating filter into service remotely, i.e. no need to drive into the office. These features will be important to controlling operational costs.

NetFlow

To plagiarize Wikipedia, “NetFlow is a feature introduced on Cisco routers to provide the ability to collect IP network traffic as it enters or exits an interface. By analyzing the data provided by NetFlow, a network administrator can determine things such as the source and destination of traffic, class of service, and the causes of congestion.”

Today, NetFlow is usually associated with metadata collection. It is a very useful way to generate meaningful information that a network administrator will be able to leverage to troubleshoot a network issue. From a monitoring perspective, it is essentially aggregated data samples from the network. This makes it different from packet data, which is a copy of the exact data in a packet. With packet data, you get copies of the actual packet, not high level flow data like you get with NetFlow. Both data formats are necessary as monitoring tools are built for only one of these specific formats, not both.

PURPOSE OF NETFLOW

The first version of NetFlow was introduced in 1996 by Cisco as a proprietary protocol. It went through some important milestones like version 5 (in 2009). It was later harmonized with the IETF Internet Protocol Flow Information Export (IPFIX) standard and versions 9 and 10 of NetFlow now support IPFIX. Version 5 and 9 are probably the most ubiquitous versions today. They are used in many network routers and switches. IPFIX, which introduced user defined flow keys, was an important improvement.

NetFlow works by having devices, like network routing equipment and monitoring equipment (e.g. network packet brokers) that are called “exporters or generators”, create the NetFlow data and forward it to other devices, called “collectors” (typically dashboards like Splunk or something), that compile the data into meaningful information.

The following parameters form the basic subset of information contained within NetFlow data:

1. Ingress interface (SNMP index)
2. Source IP address
3. Destination IP address
4. IP protocol
5. Source port for UDP and TCP (0 for other protocols)
6. Destination port for UDP and TCP, type and code for ICMP, or 0 for other protocols
7. IP Type of service

TYPICAL USE CASES

A typical monitoring setup using NetFlow consists of three main components:

- The “flow exporter” which aggregates packets into flows and exports flow records toward flow collectors
- The “flow collector” responsible for reception, storage and pre-processing of flow data received from the flow exporter
- And the analyzing application which will complete advanced flow processing to detect intrusion detection for example.

When implemented, these components allow you to enhance the following efforts:

- Reduce network and application troubleshooting time and effort
- Discover indicators of compromise that can help you improve your threat detection process
- Conduct application-level filtering
- Improve regulatory compliance initiatives
- Maximize performance monitoring efforts

CONSIDERATIONS

Here are some things to keep in mind when considering NetFlow-based monitoring solutions.

What NetFlow versions does your equipment support?

NetFlow variations occurred as the versions progressed. There are differences between version 5 and version 9 (which supports IPFIX protocol mentioned earlier). Version 9 also introduced user defined flow keys, which was a major enhancement and allows for additional NetFlow information, i.e. extensions to NetFlow to be created by vendors.

For instance, Ixia solutions like the Advanced and Threat Intelligence Processor product, using something called IxFlow (created by Ixia) that supports the following extensions to NetFlow:

- The operating system
- User browser type
- User device type (handheld, laptop…)
- Geo-location information (Country code, client and server city, etc.)
- Application information (name and ID)
- Service provider

Make sure you have the right equipment (infrastructure, packet brokers, and monitoring tools) that supports the NetFlow version you want to use.

Your monitoring needs

What problems are you trying to solve and what data do you need? For instance, do you need flow data or packet data? Do you need both? Do you need geo-location and other additional NetFlow data that are extensions to the basic protocol? These are the questions you need to ask yourself.

Network Function Virtualization (NFV)

As a Forrester report declared some years ago, "Hardware is dead — or, more precisely, it has left center stage." Companies are realizing the benefits of software-based solutions over hardware-based ones. And functions that previously could only be delivered through hardware, can now be achieved through software – all thanks to virtualization.

WHAT IS NFV?

Simply put, Network Function Virtualization (NFV) occurs when network functions (e.g. encryption, DNS, firewall services, etc.) once performed by hardware, are instead carried out using software on a Virtual Machine (VM).

NFV was conceived as an answer to the challenges associated with hardware-based network function solutions. Traditionally, companies in need of network function technology have been forced to have the relevant hardware installed on-site. Their network service provider performs this service and results in high costs, is time intensive, and creates an upgrade pain. But with NFV, the service provider can deliver these functions virtually. Virtual Machines can be used to run network firewalls and other capabilities remotely. This software-based solution incurs lower costs and is easier to deploy and upgrade.

When a network function is executed virtually, it is known as a Virtual Network Function (VNF).

USE CASE AND BENEFITS OF NFV

Instead of installing various pieces of network function hardware at the customer premises, the service provider simply installs a generic server on site. The server is a commercial off-the-shelf (COTS) product, making it inexpensive:

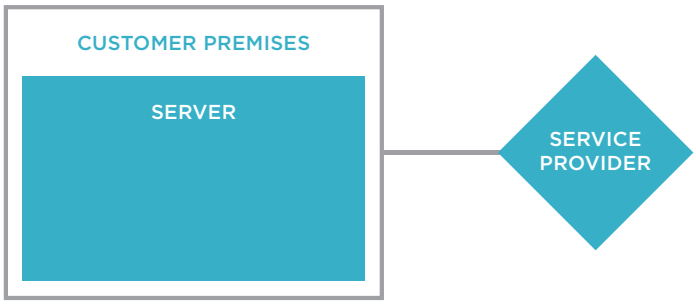

The provider then uses virtualization software (e.g. VMware) to create Virtual Machines (VMs) on the server. Each of these VMs performs the network functions previously delivered via hardware:

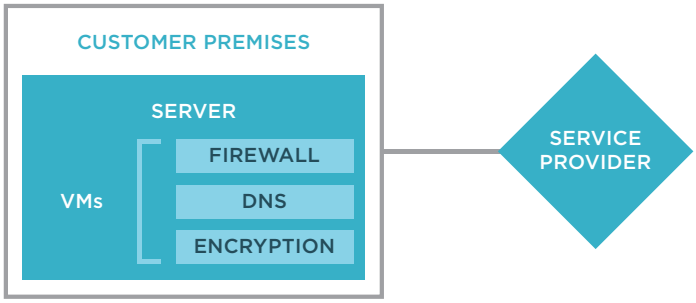

The server that runs the VMs is called a compute node. Back in the service provider's datacenter, a Virtualized Infrastructure Manager (VIM) is used to manage several compute nodes at once. The software that ultimately manages customers' network functions is also run in the service provider's datacenter:

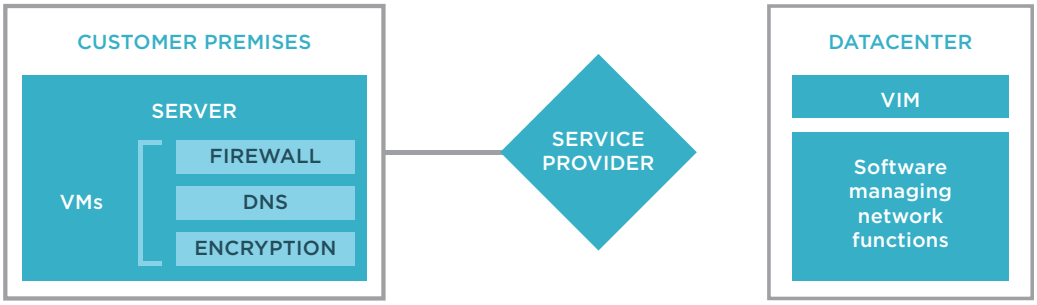

CONSIDERATIONS ASSOCIATED WITH NFV

There are several considerations, especially challenges, to consider when making an NFV purchase.

The unknown

Cloud and virtualization technology are by no means new, but they've not been around as long as hardware systems have. The risks associated with network function hardware are well known and documented. NFV has created a new, unknown environment, bringing with it fresh dangers. Moreover, it's commonly accepted that software-based solutions are not as stable as hardware-based ones. NFV is inherently more prone to security threats than physical network function solutions.

Make sure to validate your NFV solutions before rollout

Organizations can minimize the risks associated with NFV migration through testing. For instance, Ixia offers a number of testing solutions to help ensure virtualized infrastructures add value – rather than pain, complexity, and security threats.

Technical hurdles

For companies dependent on physical network appliances, the transition to virtual solutions is fraught with complexity. Just a few of the challenges which must be overcome include:

- Managing the coexistence of hardware and virtualized solutions
- Managing the deployment and coexistence of virtualized solutions from different vendors
- Managing multiple virtualized network solutions and keeping them safe from attack
- Applying automation across the board so that NFV can be scaled up with time

Data Visibility

Network visibility can be another challenge companies face when transitioning to virtualized environments. Some 80% of "east-west" traffic (traffic that occurs between and within VMs) is invisible to traditional network monitoring tools. Virtual tap visibility solutions help organizations eradicate this virtual blind-spot as they move to NFV.

MORE INFORMATION ON NFV

For more information on the NFV implementations, how testing can help mitigate these risks, and how virtual taps can remove blind spots, visit the Ixia website.

Network Packet Brokers

Staying in control of your network is a never-ending battle. Cloud computing. Virtualized networking. Growing numbers of connected devices. And increasing numbers of security threats emerge daily. Yet amid all this complexity businesses rely on their applications to be reliable, fast and secure.

Having the right monitoring and security tools has never been more vital.

Organizations often end up using a blend of tools from different vendors. But trying to provide an amazing quality of experience with this wide mix of tools gets complicated very quickly. And there's simply no room for error – since monitoring and security tools are only as good as the data they see. Network packet brokers are a practical way of addressing this problem, and re-strengthening network security.

WHAT IS A NETWORK PACKET BROKER?

A network packet broker (NPB) is a device that provides a collection of monitoring tools with access to traffic from across the network. The word "broker", or "dealer" is helpful to focus on here. The diagram below shows how an NPB receives data from a number of network links. It then acts as a "broker", or "dealer", dealing the relevant data out to the relevant monitoring tools.

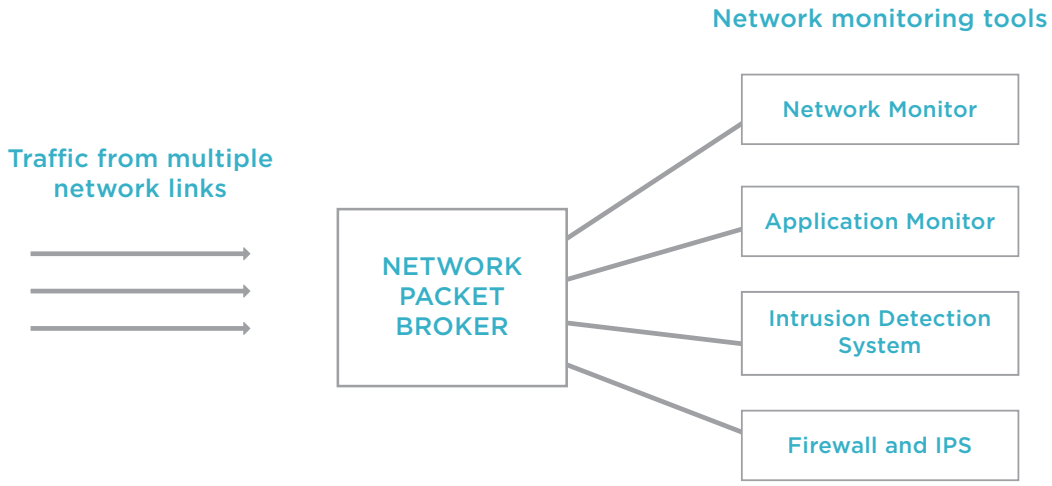

It's a middleman for network monitoring traffic. NPBs are active "brokers", or "dealers" of data, because they can be specific and targeted in the data that is supplied to each tool.

NPBs can:

- Deal data from one network link, to one tool
- Deal data from one network link, to multiple tools
- Deal data from multiple network links, to one tool
- Deal data from multiple network links, to multiple tools

Ultimately, NPBs make monitoring and security tools more effective, by giving them access to a range of data from across the entire network. Blind spots are reduced, giving tools the visibility they need to identify and tackle performance and security threats.

USE CASE AND BENEFITS

Network packet brokers aren't all made equal. NPBs are a fairly new form of technology. As a result, vendors vary in the way they design these products – there's no established NPB model.

However, there are at least four things a good NPB will do:

Safe Removal of Redundant Data

Not all traffic that flows through an NPB is useful – some data may be duplicated. To save time, and processing power, duplicate packets, and other redundant data can be removed before reaching monitoring and security tools. During this process, it's imperative that relevant original data isn't accidentally dropped. Advanced NPBs offer zero-loss advanced packet processing at full line rate. This means redundant data is carefully sifted out, while all important, original data packets are preserved and provided to your tools.

Application Intelligence and Filtering

Managing the network means knowing what's on the network. Large networks can have hundreds of applications running, especially with the growth in BYOD. Intelligent NPBs can identify the applications in use on the network and provide that intelligence to any of your tools. Many tools in use only need to monitor or inspect specific types of applications. Intelligent NPBs can easily "broker" or "deal" traffic out to monitoring and security tools by application flow. This makes your monitoring and security tools much more efficient, and it makes life much easier for administrators.

SSL Decryption

Secure Socket Layer (SSL) encryption is the standard technology used to send private information. While it helps protect sensitive data, it also comes with network security risks. SSL hides sensitive data – but it can end up encrypting and hiding malicious cyber threats too. For network safety, organizations must decrypt and examine SSL traffic. But decryption takes up valuable processing power. If decryption is left to security tools to perform, time is wasted unraveling code, rather than scanning, and keeping your network safe from threats. Intelligent NPBs can perform SSL decryption, passing on the decrypted data to your monitoring tools. This gives them the ability to see all of the traffic, and the freedom to get on with protecting your network.

Data Masking

One drawback to SSL decryption is it makes all data visible to anyone that has access to your monitoring tools. Some of this unencrypted data may be quite sensitive and protected by regulatory requirements. Data like personally identifiable information (PII), or credit card information must be protected and not exposed to unauthorized individuals. Thus, advanced NPBs can mask unencrypted sensitive data that should not be, and does not need to be exposed to monitoring and security tools or their administrators. This data masking can be a critical NPB feature that makes monitoring activities safe.

CONSIDERATIONS

To eliminate future headaches and problems, make you're your visibility architecture includes the above features and more for your network packet broker solutions. Setup and configuration of these intelligent features are made easy with intuitive, drag-and-drop graphical user interfaces (GUI) that enable you to get on with the business of securing your network.

Network Taps

It takes 20/20 vision to keep networks up and running. From data packets to devices, you need to be able to see everything on the network clearly. It's simply what it takes to protect your network, and ensure peak performance. Anything less than 100% visibility can leave you exposed to network failure. But how can a true view of the network be achieved? Network taps offer a low-cost, complexity-free solution.

WHAT IS A NETWORK TAP?

A network tap is a way of monitoring the data flowing across a network.

Picture a network where information flows between devices A and B:

A network tap can be placed between the devices. The tap lets traffic through unhindered – but it also creates a copy of the data that passes through:

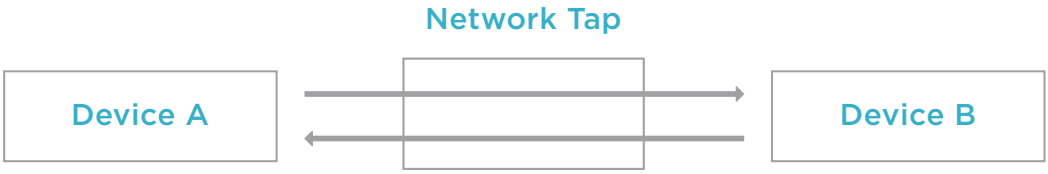

When a monitoring device is connected to the tap, it is able to view this duplicated network traffic:

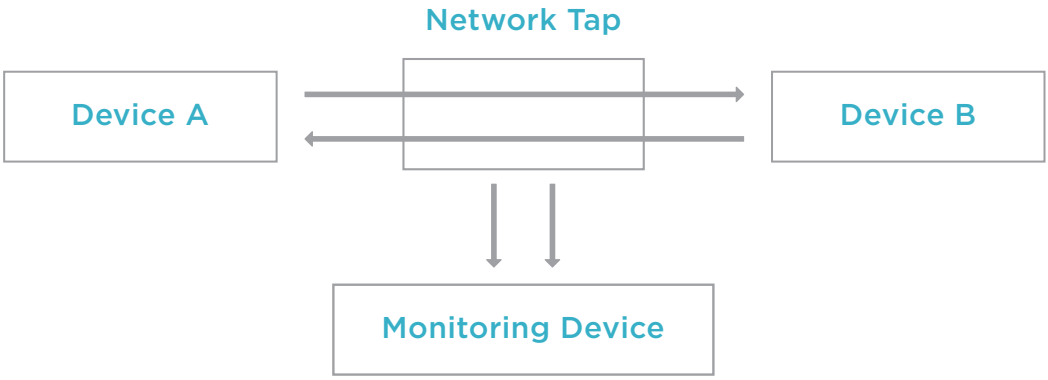

TYPICAL USE CASE AND BENEFITS

What's so great about network taps? With other monitoring methods, network visibility comes at the expense of network performance. Seeing your network clearly can mean sacrificing speed, or interrupting traffic flow. But network taps give maximum visibility with minimal disruption, and can be much less expensive in the long run than their SPAN port (switch port analyzer) alternatives.

Once taps are in place, it's easy to connect monitoring devices to the network without impacting performance. Devices that help keep networks safe (e.g. intrusion detection / prevention systems) and smooth running (e.g. protocol analyzers, RMON probes), can be deployed quickly and painlessly. And network taps "fail open" too. Even if a tap loses power, or stops working, network traffic will continue to pass through it unaffected.

Don't SPAN ports do the job too? SPAN ports, also known as port mirroring, is another method of monitoring network data. But SPAN ports aren't nearly as effective as network taps:

- When SPAN ports reach capacity they stop capturing full data
- SPAN ports can introduce delays to the network
- SPAN ports miss or corrupt data packets, and may not capture errors
- SPAN ports are vulnerable to attack

In contrast:

- Taps receive all network traffic – including errors
- Taps don't cause network delays
- Taps don't change the content or structure of network data
- Taps don't have a network address and cannot be hacked

Watch our Taps vs. SPAN video, or download our Taps vs. SPAN brief for more on the benefits of taps over SPAN.

CONSIDERATIONS FOR TAPS

As communications evolve, the need to monitor traffic on networks carrying VoIP, video-conferencing, and security applications is growing. Taps are mission-critical for network security and performance. And they outperform SPAN ports in a number of ways. But not all network taps are created equal. Make sure your taps capture all network traffic without introducing bottlenecks or points of failure. As an example, Ixia copper taps come with backup power sources that kick in if the primary source fails. Good optical taps should give you choice the choice of a variety of split ratios (from 50/50 to 90/10) so that you can pick the right tap or for your light budget. Make sure the tap you pick fits your needs.

Service Chaining

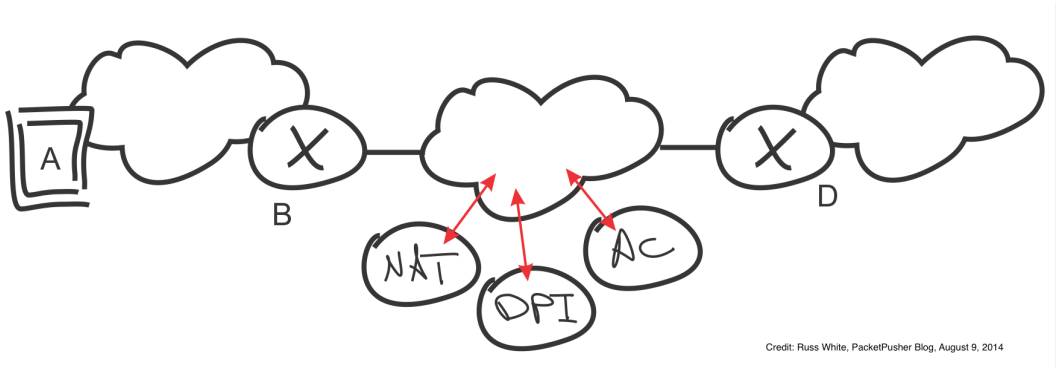

Service chaining is a powerful tool for automating the movement of data packets and it's getting a lot of attention these days as a way to improve the quality and speed of application delivery.

PURPOSE OF SERVICE CHAINING

A 'service chain' is a set of network services which are performed in a specific order and 'service chaining' refers to steering the traffic through such a chain. It's like a recipe where actions are performed in a pre-ordained order. Services can be performed in parallel or in serial, depending on the situation. The chain can be implemented by cabling individual devices together or, increasingly, by using software provisioning to control the flow of data through the selected services. Monitoring tools that are linked together in this way are sometimes referred to as a daisy-chain.

The use of service chains is linked to the automation of functions that have been either embedded in single-purpose hardware devices, dictated by physical topologies, or performed manually--which are increasingly perceived as too costly and inflexible in our fast-moving digital economy.

TYPICAL SERVICE CHAINING USE CASES AND BENEFITS

Service chaining is one of several approaches that make it possible to centrally manage and direct the operation of IT resources, to increase efficiency and time-to-market, as well as decrease costs.

Real-Time Network Monitoring

With real-time monitoring, you need to keep traffic moving quickly and your security tools working efficiently. Chaining tools together allows to you to pass only the suspicious traffic to additional tools for deeper inspection or to a honeypot to be quarantined. Packets without anomalies are moved along quickly, to maintain maximum response time. A common example is the use of a Security Information and Event Management (SIEM) solution to filter out suspicious traffic for further analysis by other tools in the daisy-chain. Traffic without exception is quickly sent back through the network to support the fastest possible response time.

Out-of-Band Monitoring

Out-of-band monitoring tools can be chained for similar reasons. An example would be to take the result of deep packet inspection provided by an Ixia network packet broker and send the application-specific information on to the best tool for analyzing a given packet type. Meta data can also be added to the packets to let tools farther in the chain know more about the origin or destination of the traffic.

Value Added Traffic Management

Service chaining is also common when administrators must enable multiple resources or processes to be used. Examples are to enforce policies, perform QoS monitoring, to gather real-time analytics for traffic flow adjustments, are enforced to ensure quality of service.

Service Management

The concept of service chaining plays a strong role in helping carriers provide services to end users with speed and accuracy or helping providers deliver a service with an excellent experience. One example is the chain of special-purpose platforms that video packets must pass through before delivery to the end customer, beginning with video optimization, then transparent caching, then (optional) parental controls, and finally a WAP gateway. These services are linked or chained together so that tasks necessary for all of these services do not have to be performed multiple times. Details about each user—such as their device, location, or whether they are subject to parental control—are also used to dynamically steer traffic through the necessary services.

ADVANTAGES OF SERVICE CHAINING

Enable Network Function Virtualization (NFV)

Once upon a time, specialized network appliances ruled the data center and in many places they still do. When you consider their purpose, however, you can identify multiple functions taking place inside each appliance. For instance, a firewall might perform network address translation, deep packet inspection, and access control. The hardware appliance was designed to perform these functions at wire speed. But in recent years, many of the functions once performed by expensive hardware appliances are being redesigned as software functions that can be run on any generic and low-cost CPU. This process is called network function virtualization and the goal is to achieve the same results as the appliance, but at greater efficiency and less cost.

Reduce Latency

In order to get acceptable performance in a virtualized environment however, services that run in software on a generic CPU must be chained together, to accelerate total processing speed or latency. Any time services are grouped together in a way that forces processing to proceed from step one to step two, latency can be reduced and speed accelerated.

Reduce Redundant Inspections

Without the ability to chain together certain functions, a particular packet may need to pass through a particular service more than once to meet the qualifications for other types of inspection tools. For instance, in the case of security monitoring, SSL traffic can pass through a powerful decryption tool and the exposed content can be sent through a series of additional inspection tools. This avoids the need to send the traffic through decryption for each tool, which would increase latency and multiply the cycles being consumed on the decryption tool. A more efficient and more cost-effective result is achieved by sending decrypted traffic through multiple tools before passing it through to the trusted network.

Apply Consistent Policies

Pre-set service chains help ensure that actions are taken in a specific sequence and nothing is overlooked. This reduces errors and increases the chance that abnormalities will be identified in time to prevent damage to an organization's data or other resources.

Increase Flexibility

The ability to define service chains dynamically, based on the user, device, location, service level, or other characteristic is a powerful capability in the fast-moving digital economy. Well-defined rules and policies can help decrease the time to deliver a service and increase the quality of the user experience.

CONSIDERATIONS

Service chaining is a useful concept that can help you organize operational tasks into more manageable groups. As programmability becomes the norm in network management, organizations will find more ways to use service chaining to increase network visibility, improve security monitoring, and increase the speed and quality of applications.

SSL Decryption

Coded, disguised forms of communication (also known as ciphers) are nothing new. For centuries, people have sent messages designed to be unreadable, if intercepted or received by hostile forces. The enigma machines used by Germany during WW2 are a famous example of this.

Today, SSL (secure sockets layer) encryption and decryption are the means by which sensitive data is safely transmitted – and protected from prying eyes – over the internet and across networks.

On a basic level, SSL encryption occurs when sensitive data is transformed into an unintelligible, unreadable "ciphertext". SSL decryption occurs when this "ciphertext" is returned to its original format. To do this, a key, which gives instructions on how to decode the encrypted message, is required.

NETWORK SECURITY AND SSL DECRYPTION USE CASES

We live in an age where the stakes are high for both individuals and organizations that fall victim to data theft. So it's for good reason that SSL encryption has soared (and continues to soar) in popularity. According to Mozilla, half of all internet traffic is now encrypted. SSL encryption is a powerful weapon in the battle for data security, but its greatest strength, is also its greatest weakness.

Encryption hides at-risk data. But it can also hide other, less innocuous things too. Cybercriminals can take advantage of SSL encryption, camouflaging malware and other undesirables in encrypted data, so that they're able to sneak into, and around company networks undetected.

Since most network tools can't inspect SSL encrypted data, it's necessary that data is first decrypted, and then inspected.

CONSIDERATIONS FOR SSL DATA MONITORING

So SSL decryption is vital to network security, yet it presents a number of challenges. These include:

Firewall Strain

Many organizations use their Next Generation Firewalls (NGFWs), (which typically come with decryption capabilities) as the main point of SSL decryption on the network. But, as data encryption continues to grow in popularity, firewalls will experience more and more strain. SSL decryption takes up valuable processing power, making the firewall less effective. This makes it easy for the firewall to become a bottleneck on the network, holding up other processes and stalling productivity.

Internal Threats

Following on from the above, firewalls often sit at the edge of the network, only decoding encrypted data that comes from outside. This means that encrypted communications that occur internally (between servers / clients) remain uninspected. This is a huge risk, since internal communications may comprise some 80% of encrypted network data. If malware does somehow make it into the network, SSL encryption within the network will help camouflage it.

Inefficiency and Overloading

Some network monitoring tools (aside from NGFWs) come with SSL decryption capabilities too. This isn't an ideal solution however. As with firewalls, enabling SSL decryption on these tools can debilitate performance. Furthermore, requiring each tool to decrypt its own data is inefficient. It means multiple, siloed tools, performing the same decryption process, on the same set of data. This is a waste of resources. Why have several appliances repeating the same task when one tool could decrypt the data, and then push it out to them all?

Data Compliance

HIPAA, PCI, and organizational best-practice policies mean personally identifiable information (PII) must be handled with care. When SSL decryption occurs, sensitive, personal data is exposed. Without the right safeguards in place, anyone with access to network monitoring tools becomes privy to this sensitive information. Data masking should be used to protect this data.

MORE INFORMATION ON DECRYPTING MONITORING DATA

Companies can overcome the challenges associated with SSL decryption by using a network packet broker (NPB) with decryption capabilities. Network packet brokers act as a central recipient of and "dealer" of data in the network. By decrypting the data at this point, before sending it out to network appliances (e.g. Intrusion Detection Systems (IDS), Intrusion Prevention Systems (IPS), Firewalls). The NPB removes the processing burden from these tools, boosting overall efficiency and productivity. Unlike firewalls, NPBs are also able to decode encrypted data that arises from within the network.

A good NPB solution comes with built-in data masking capabilities. Designed with usability in mind, a point-and-click user interface takes the pain out of data compliance in this area.

Find out how NPBs work here. Alternatively, read this article on Ixia's data masking capabilities.

Traffic Filtering

How do you find a needle in a haystack? In MythBusters Episode 23, Jamie and Adam each ended up destroying the hay in order to find the needle. But this is not an option for network and security administrators. Monitoring and securing modern networks requires finding "the needle" without destroying the network, or even the network traffic. Very sophisticated and automated analytics tools make this possible.

Specialized tools like:

- Network performance monitoring and diagnostics (NPMD)
- Application performance monitoring (APM)
- Next-generation firewalls (NGFW)
- Intrusion detection systems (IDS)
- Intrusion prevention systems (IPS)

It is simply not possible to manually monitor and secure our networks without these automated tools. But these tools are expensive. How do we get the most from our monitoring and security investments?

This is when most companies start investigating the use of network packet brokers (NPB). Intelligent NPBs have a whole host of features, such as those listed below.

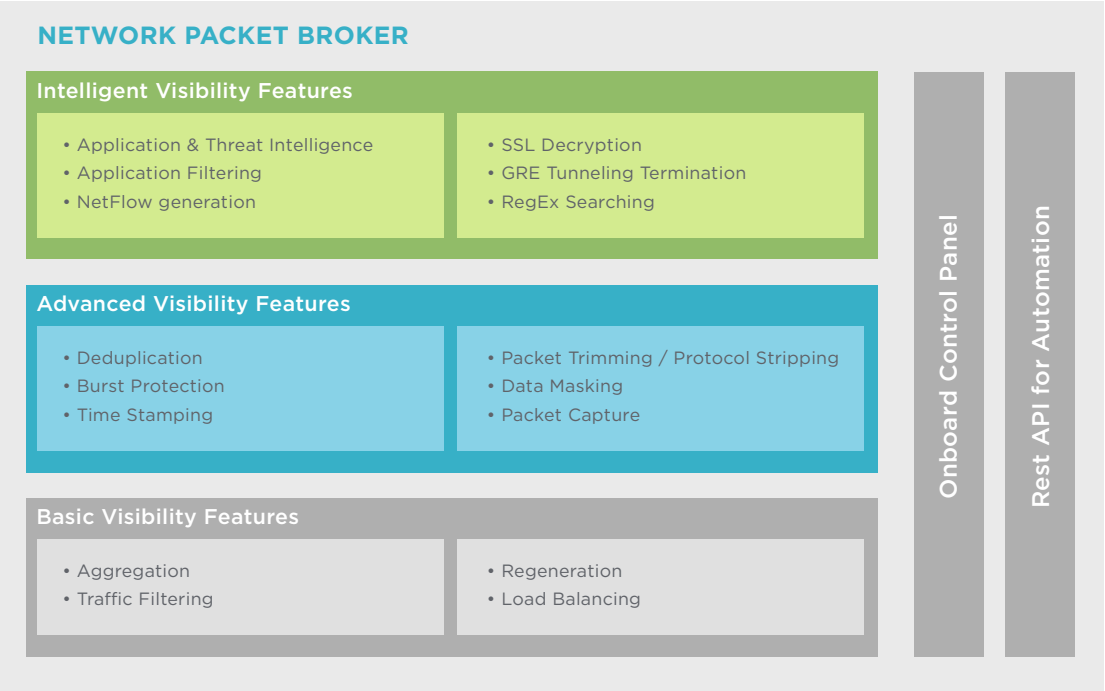

Traffic filtering and application filtering (traffic filtering by Layer 7 application flows) are features that directly help get more visibility and security from less monitoring and security tools capacity.

PURPOSE OF TRAFFIC FILTERING

Seeing more with less. Securing more with less. Traffic filtering is all about delivering on these promises. Traffic filtering allows the user to define:

- Specific traffic I want to see excluding all other traffic (pass by criteria)
- Specific traffic I do not want to see accepting all other traffic (deny by criteria)

Each approach is designed to limit the amount of data sent to monitoring and security tools making the individual tools much more efficient. After all, it is easier to find a needle in half a haystack than in a whole haystack.

Traffic filtering allows us to reduce the amount of tool capacity needed. Or we can use the saved budget dollars to purchase additional tools that we might not have had the funds for otherwise.

In short, the purpose of filtering traffic for our monitoring and security tools is to:

- Make the tools operate more efficiently.
- Make optimal use of the tool capacity.
- Reduce our overall investment in individual tools sets.
- Allow us to purchase additional tools.

TYPICAL USE CASES

Here are some real life use cases where traffic filtering is beneficial.

- **Reducing Security Tool Costs** – The University of Texas at Austin recently deployed intelligent NPB's in their network. Using Ixia's application filtering, they were easily able to send some traffic, like student Netflix movies and streaming music around their IDSs. This simple step reduced the load on their security tools by 20-30% and delivered a 100% return on investment.
- **Improving Voice and Video Monitoring** – Citrix unified communication services are a critical productivity application for many organizations. Monitoring quality of experience can require analyzing SIP based and PSTN originated call data. However, the VoIP call data and PSTN call data are analyzed on different tools. Traffic filtering easily sends only the relevant traffic to each tool for analysis.
- **Filtering Encrypted Traffic for Decryption** – According to Gartner research, secure sockets layer (SSL) traffic is a significant portion of all outbound Web traffic and is increasing. It represents on average 15-25% of total Web traffic, with strong variations based on vertical market.1 Unfortunately, it prevents monitoring and security tools from inspecting the traffic. So SSL decryption is required. With an intelligent NPB, companies can use application filtering to identify SSL traffic and send only this traffic to SSL decryption tools or internal SSL decryption capabilities. Here application filtering saves as much as 80% of the capacity of SSL decryption tools.

- **Expediting "On-the-Fly" Troubleshooting** – Reducing trouble resolution times is a critical metric for IT organizations. Filtering traffic "on the fly" for forensics tools or built-in packet capture is an important troubleshooting feature on NPBs that helps significantly speed trouble isolation and reduce resolution times. In fact, customers have experienced as much as 80% reduction in troubleshooting times.

CONSIDERATIONS WHEN RESEARCHING NETWORK PACKET BROKERS

Traffic filtering can be one of the most complex operations performed on any NPB. So it is critical to know what to look for when evaluating these tools. Below are some important NPB traffic filtering selection criteria.

Layer 7 Application Filtering

Being able to easily route application flows is critical for network visibility and security. Many NPBs only route traffic on Layer 2-4 protocols.\ Consider an NPB with deep packet inspection and application intelligence that can filter each application flow or even make RegEx filtering simple.

Operational Ease of Use

Configuring traffic filters can be extremely complex. Consider a NPB that automates the entire traffic filtering process, which eliminates all that complexity. Do not get forced into a solution that requires your team to manually deal with the traffic filtering complexity.

No Dropped Packets

Traffic filtering can be computationally difficult if not implemented well in an NPB. Consider only NPB's that maintain the complete packet stream when filtering is enabled. An NPB should NEVER drop packets, not even when users are making multiple simultaneous traffic filtering changes to the NPB configuration.

Simultaneous Usage

NPBs are typically used by multiple teams within an IT organization. Consider an NPB that supports simultaneous usage by multiple team members without causing any traffic filtering issues or configuration errors. Without this capability, teams can bump heads when emergencies arise.

CONSIDERATIONS

Finding a needle in a haystack is difficult. Network packet broker traffic filtering capabilities help monitoring and security tools do the job much more efficiently. But choose wisely. Not all traffic filtering capabilities on NPBs are the same.

Sources:

[1] Gartner, "Security Leaders Must Address Threats From Rising SSL Traffic," Gartner,December 9, 2013, refreshed January 8, 2015

Virtual Taps

A virtual tap is a software-based solution that captures a copy of the data flowing between virtual machines (VMs). Virtual taps provide clear visibility into inter- and intra-VM traffic (also known as east-west traffic). They're able to copy (also called mirror) VM data, filter the mirrored data, and then send the mirrored traffic of interest to physical or virtual monitoring tools.

PURPOSE OF VIRTUAL TAPS

So, what's the overarching benefit that they provide? Virtual taps eliminate virtual network blind spots and enable IT to analyze critical data for security threats and performance issues. For example, virtual traffic may be sent to inline security tools such as Intrusion Prevention Systems (IPS), or out-of-band security tools, or even sent to performance monitoring tools.

Research we conducted last year revealed that 80% of modern enterprises consider server virtualization a strategic priority. In addition, two in three companies run critical applications on virtual servers. Not only is virtual computing on the rise – it's also being used to perform crucial business tasks. But to see all of the data flowing across virtual environments, you need a virtual tap.

TYPICAL USE CASES

With organizations increasingly dependent on virtual computing, it's essential that virtual environments remain smooth-running and secure. But virtual networks are particularly vulnerable to performance challenges. And they're a natural target for security threats too.

Here are some real life use cases where virtual taps are beneficial.

- **Strengthening security defenses** – Virtual taps are the best defense against costly cyber threats in virtual environments. They enable the oversight ability to detect security risks. Malware variants like Crisis have been optimized to function in virtual environments. Without visibility into your east-west traffic, how do you know you haven't already been compromised? What would alert you to this fact?
- **Reducing performance issues** – Virtual taps give you access to performance data in your virtual data center. Network and data center failures can be costly, especially when unplanned. Virtual taps give you the visibility you need to perform trending analysis to avoid potential component problems and tackle operational issues.

- **Consolidation of regulatory compliance initiatives** – Many organizations need visibility within virtual environments in order to comply with service level agreements (SLAs) and other industry regulations (e.g. HIPAA in healthcare, PCI-DSS for financial card transactions, SOX in the enterprise). By capturing data from your virtual data center and exporting it to your existing compliance tools so that it can be combined with data from your physical data center, you now have complete network visibility and can demonstrate that visibility as part of any compliance audit.

CONSIDERATIONS WHEN RESEARCHING VIRTUAL TAPS

When considering virtual taps, there are several items to investigate. Here is a short list of common items:

Multiple hypervisor support

You will want a virtual tap that supports the most common hypervisor types, like VMware, Hyper-V and KVM. Even if you only have one VM type in your network, multi-hypervisor support gives you flexibility down the road for additions and change of direction in your virtual data center.

Single pane of glass to see your virtual taps

Once you have all of your virtual taps installed, you want to be able to see them in one consolidated view. This includes virtual taps installed in different VM environments. You will want to see them all from a single pane of glass so that you can more efficiently and cost-effectively monitor your network. Siloed views of virtual taps not only create irritation and loss of productivity, they can obscure oversights in your monitoring plans and cause confusion.

Performance impacts

A third consideration is around performance. You want a virtual tap solution that does not create any significant performance issues for your network. This includes not adding significant load to the CPU or the VM and also not overloading the LAN (due to the creation of the mirrored data). The virtual tap must be able to filter the mirrored data before it is sent across the LAN. Otherwise, you'll heavily load your network with the extra 50% to 100% of mirrored data.

Visibility Architectures

Visibility is defined by Webster as the “capability of being readily noticed” or “the degree of clearness”. For network or application visibility, we are talking about removing blind spots that are hiding the ability to readily see (or quantify) the performance of the network and/or the applications running over the network. This visibility is what enables IT to quickly isolate security threats and resolve performance issues; ultimately ensuring the best possible end user experience.

Another way to think about this is that visibility is what allows IT to control and optimize the network along with applications and IT services. This is why network, application, and security visibility are absolutely vital for any IT organization to accomplish their job! Without visibility, IT can only operate reactively to problems and will never be truly effective at eliminating those problems.

A Visibility Architecture then is the end-to-end infrastructure which enables physical and virtual network, application, and security visibility. While it is possible to piecemeal visibility components together as you solve one problem after another, this won’t give you a cohesive strategy. That practice would only lead to unnecessary complexity and far higher costs. The basis of a visibility architecture starts with creating a plan. Instead of just adding components as you need them at sporadic intervals (i.e. crisis points), step back and take a larger view of where you are and what you want to achieve. This one simple act will save you time, money and energy in the long run.

A proper visibility architecture addresses the strategic end-to-end monitoring goals of the network, whether they are physical, virtual, out-of-band, or inline security visibility. Once you combine the security architecture with the visibility architecture, you will equip yourself with the necessary tools to properly visualize and diagnose the problems on your network.

PURPOSE OF VISIBILITY ARCHITECTURES

A visibility architecture typically yields immediate benefits such as the following:

- Eliminating blind spots
- Controlling costs while maximizing ROI
- Simplifying control

First, you want to eliminate blind spots, i.e. the hidden areas of your network. Every network has some. By designing an architecture, you have a full array of solutions for both physical and virtual deployments that can be leveraged so that network operators don't have to make compromises when it comes to visibility. For an extensive list of blind spots, see this blog.

The starting point of your architecture is to make sure you have proper access to the data you need. This typically involves using taps, virtual taps, and bypass switches to access data from relevant segments of your network. This removes the bottle neck caused by limited access points (like SPAN ports). However, SPANs can still be used, if necessary.

Next, you'll want to have a filtering component to maximize the flow of relevant information to your monitoring tools. Enterprises can maximize their monitoring investment by utilizing powerful network packet brokers (NPBs). These devices give greater control to network operators and enable the ability to extend the life of existing network, application, and security tools; even as you migrate to higher speed 10GE, 40GE, or 100GE networks. NPBs are responsible for data aggregation, filtering, deduplication, and load balancing of Layer 2 through 4 (of the OSI model) packet data. These features ensure the tools get the data they need without being overwhelmed.

The next set of capabilities is the application intelligence layer. This functionality allows filtering and analysis further up the OSI stack at the application layer, i.e. Layer 7. These capabilities give you quick access to information about your network and help to maximize the efficiency of your tools. This is only available in certain NPBs. Depending upon your needs, this feature can be quite useful as you can collect the following information: the types of applications running on your network, the bandwidth each application is consuming, the geolocation of application usage, device types and browsers in use on your network, and the ability to filter data to monitoring tools based upon the application type. You can also perform SSL decryption at this layer.

The final layer is made up of your security and monitoring tools. These devices perform the analysis function on the security and monitoring data. They are typically special purpose tools (e.g. IPS, firewall, sniffer, APM, etc.) that are designed to analyze specific data. The output from these tools is typically used by network engineers to make their decisions.

Typical Use Cases

When all components of a visibility architecture are combined, they eliminate the blind spots within your network that are harboring potential application performance and security issues. Here are some real-life use cases that show off the benefits of a visibility architecture.

- **Strengthening of security defenses** – A primary reason for a visibility architecture is because if your network is attacked, or breached, how will you know? A DDoS attack will usually impact website performance. But other than that, how will you "see" a security attack? This is actually a common problem. The 2015 Trustwave Global Security Report stated that 81% of compromised victims did not detect the breach themselves—they had no idea this had happened. The report also went on to say that the median number of days from initial intrusion to detection was 86 days. So, most companies never detected the breach on their own (they had to be told by law enforcement, a supplier, customer, or someone else) and it took almost 3 months after the breach for that someone else to notify them.
- **Acceleration of Mean time To Repair** – Another example of a visibility architecture benefit is faster remediation of security breaches and network problem. In regards to security problems, if you can't see the threat, how are you going to respond to it? For network problems, where should you start your troubleshooting efforts? A visibility architecture gives you a coherent way and access to the data you need to triangulate on problem spots as fast as possible. Some Ixia customers have seen an up to 80% reduction in their mean time to repair performance due to implementing a proper visibility architecture.
- **Prevention of network issues and problems** – Prevention is always a good aspiration. Almost all of us have grown up on the phrase, "an ounce of prevention is worth a pound of cure." With proper visibility into your network, you can capture data the data you need to prevent costly outages. For instance, network data can tell you applications or network segments are running slowly. You can even run proactive monitoring solutions to test network segments and applications to check that they are working normally or see what kinds of problems they are having. Application intelligence can also help in this area.
- **Optimization of your network** – The final goal of a visibility architecture is to be able to capture data at regular intervals so that you can characterize your network and understand where and when you might have issues. This allows you to be even more proactive. In addition, it gives the data you need to better dimension your network equipment, optimize traffic routes, and maximize your capital expenditures (CAPEX).

CONSIDERATIONS WHEN RESEARCHING VISIBILITY ARCHITECTURES

When considering visibility architectures, there are several items to investigate. Here is a short list of common items:

Flexibility, i.e. choice

You will want, and need, options. This includes the flexibility to deploy inline and out-of-band visibility solutions. It also includes the ability to monitor your physical and virtual data center traffic. Application Intelligence is another area to look for. While you may not want to engage in all of these activities right away, you should look for a solution that allows you to add the pieces you want, when you want, without a forklift upgrade.

Ease of Use

This will be a critical component that will heavily influence your total cost of ownership (TCO). Look for a solution that uses a drag and drop GUI interface. A command line interface will take you 10 times (or more) longer than a drag and drop interface to configure filters. The management system should also be able to handle everything—from global element management, to policy and configuration management, to data center automation and orchestration management. Engineering flexible management for network components will be a determining factor in how well your network scales.

Technology

A third consideration is around the technology. Buyer beware applies to this market place (just like others you are used to). While Vendor products may sound the same, they usually aren't. In general, a strong consideration should be to purchase NPBs that run at line rate under all conditions. Only a very few NPBs do this. Anything less adds delay to your monitoring effort.

For inline solutions, this line rate will be absolutely critical. You will also want failover technology that is as fast as possible for inline solutions. It also suggested to use dedicated bypass switches, instead of bypass switches built into monitoring tools. This will maximize your fail-over capabilities and minimize loss of data on your network.

Data access

Data access is another area of concern. Consider using Taps instead of SPAN ports for your data access technology. Taps are superior to SPANs for several reasons, see this analysis here. One key difference is that SPANs provide summarized data (instead of a complete copy of all data) that can often be missing key data you need for proper problem resolution. Another area to investigate is whether your tools need packet data or NetFlow data. One last thing to consider is if your tools need additional data from application intelligence functions to further improve their performance.

MORE INFORMATION ON VISIBILITY ARCHITECTURES

Visibility architectures can, and should, integrate with your business initiatives. If implemented correctly, they can seamlessly integrate into your existing network management and orchestration systems. They can also extend data center automation or application performance monitoring initiatives. Advanced visibility architectures can also take advantage of the power of automation. For example, your tools could automatically, without manual intervention, request specific types of traffic when it detects a security issue. And if a tool goes down, load balancing can automatically compensate for this by sending traffic to the remaining tools until a replacement tool can be deployed.

So what does a successful visibility architecture look like? Check out the material available here.