VISIBILITY EQUIPMENT BUYERS GUIDE

The Technical and Financial Impact of Ease of Use on Network Visibility Solutions

Written by Tim *The OldCommGuy*™ O'Neill

SECOND EDITION

TABLE OF CONTENTS

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

Foreword

This article is a comprehensive review of the strengths of GUI-based packet broker configuration and the use of taps for easy data collection. Although the strengths of these two data access methods may seem implied, it is enlightening to see their true ROI represented in numbers. GUI-based configuration is not just about improving speed but accuracy. In this article Tim O'Neill thoroughly reasons out why this methodology saves companies money and helps them optimize network engineer effectiveness. It is definitely worth a read if you are considering stepping into the packet broker arena, with a thorough explanation of the current offerings on the market, as well as the strengths of each one.

Chris Greer
Packet Pioneer

Chris Greer is a Network Analyst for Packet Pioneer LLC and a Certified Wireshark Network Analyst.

About Packet Pioneer https://packetpioneer.com/

Packet Pioneer LLC is a network analysis and troubleshooting company. In addition to diagnosing and resolving network and application issues, we also write whitepapers and technical blogs, as well as conduct training on network operation and troubleshooting using a variety of tools.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

EXECUTIVE SUMMARY

Total cost of ownership (TCO) is important to any monitoring decision. But TCO is more than just the typical security and monitoring tool purchase costs. It includes additional value-add components like Taps and network packet brokers (NPBs). More importantly, it includes the short and long term “cost of use” for the monitoring equipment as well.

So, how can you improve the short term and long-term operating costs for your monitoring solution? There are two easy steps. The first step is to update your processes to take advantage of the best technology. This means using Taps instead of SPANs to access the proper monitoring data. Better data reduces your troubleshooting and forensic analysis costs, as well as the cost due to missed security threats. In addition, you’ll want to add a network packet broker to optimize your filtering methodology and related filter programming costs.

The second step is to optimize the ease of use benefit. Ease of use includes installation, training, and day to day programming complexity. Simple choices, like using a graphical user interface (GUI), can cut your long-term operating costs by 75% or more. This is because a GUI creates higher productivity, while facilitating a lower cognitive load. By combining both steps, you can effectively reduce your TCO and reuse the extra money to solve additional needs that you have.

After looking at this criteria, several network monitoring solutions were compared. The solution from Keysight was found to be the best due to the power and simplicity of the NPB filtering engine and the intuitive capabilities of the GUI interface. Both of these components combine to create a lower total cost of ownership for network visibility (monitoring) solutions.

Initial and ongoing training costs should be a primary consideration in choosing network visibility solutions. When calculating the ROI for network monitoring solutions, the following costs needs to be factored in:

- Minimizing or avoiding network outages by early recognition of issues before failure
- Reduction of support calls to reduce operational employee production downtime
- Reducing the time to recognize and to fix issues – data leaks, attacks, etc.
- Successful management and proof of meeting corporate SLAs
- Network downtime due to lack of visibility
- Salary/staff time costs when off training plus loss of network visibility during training
- Staff frustration and overtime to find and mitigate issues that were missed due to poor visibility

In addition, it should be noted that typical costs to hire a Network Manager or Network Security Manager are very high. Losing just one of these employees due to the frustration of not having the tools required to be successful is a staggering amount of money and loss of time. The average seasoned manager will cost about $150K/year minimum, plus relocation costs, recruitment costs and the time to get the manager up and running on your particular network. Losing this kind of personnel can also be a major security risk.

The fundamental consideration is that while many vendors talk about the technical capabilities of a product, no technology is valuable unless the product can be used easily with repeatability and confidence without extensive support costs.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

CHAPTER 1: Understanding Long Term Equipment Costs

Network equipment usability (i.e. ease of use without constant training) and its contributing value to network success and security are the two most important factors to be considered when it comes to the total cost of ownership of network visibility purchases. Decisions about the quality, longevity, and ease of use of network components affect how successful an IT Team will be. This includes your security architecture, troubleshooting efforts, and network performance optimization—basically the real corporate value of your network. Your network should be considered a tangible value-add component in the overall corporate net worth.

While solutions that technically solve your problem(s) are fundamentally important, long-term usability will dictate if the solution is actually viable. For instance, any solution that technically solves your needs is a good starting point. The real question is can you sustain that level of effort, or do operating expenses (OPEX) accrue quickly for that solution?

For visibility architectures, there are four fundamental points of consideration that dictate the effectiveness of the solution:

- The security and monitoring tools purchased
- The capture of good monitoring data
- Proper data filtering
- Operational simplicity

The first consideration is the type of security and monitoring tools purchased. The network and business needs typically determine the type(s) of tools that you will need. So, there is minimal optimization capability here—from a usability perspective. This is not the case for three other considerations.

Good data collection is dependent upon the network access device and where it is located. For instance, you may want to collect data at the ingress and egress to the network and process live data with security and monitoring tools. This scenario often uses a bypass switch that is placed inline with the flow of traffic. Other situations, like network troubleshooting and forensic analysis, involve significant time delays for the data analysis so once the data is captured by a traditional Tap or SPAN, the monitoring data is siphoned off and is no longer part of the main data propagation stream. In a third situation, virtual Taps can be used to collect data from virtual data centers and cloud solutions for analysis.

Once the data is captured, it needs to be properly filtered so that the right data is sent to the right security or monitoring tool(s). This is best accomplished by a network packet broker that can regenerate, filter, load balance, and deduplicate data at line speeds, whether the network speed is 1 Gbps or 100 Gbps.

Some people may not be familiar with a network packet broker. A packet broker is really a very sophisticated filtering device. The NPB was designed for network and security managers to gain visibility access to general and specific network events and components. The main value proposition includes: ease of filter programming, programming test and verification, repeatability, and no data losses during the actions programmed and performed.

When some of the first NPB's came out, they included a simple command line interface (CLI)-programmed SPAN chip from vendors like Broadcom. The first real GUI with full bandwidth capability was the Anue (now Keysight) NPB.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

The original NPB design was created so that it would be able to filter the requisite data and feed that data to slower, but more analytical and focused devices (like performance monitoring devices, security monitors, etc.) from larger bandwidth connections. For instance, Wireshark on a PC is really only appropriate for limited 1 Gbps traffic. But what if you have a 10 Gbps or higher line rate? With a packet broker, you can send filtered data to the Wireshark or other tool, thus lowering the traffic speed and bandwidth to the different devices for a filtered view of your network. This allows slower devices to still be effective at network and security management as network characteristics change.

The last area that can help reduce the solution TCO is the simplicity with which it can be implemented and maintained. Basically, how long does it take to install and initially configure the NPB, as well as the time and effort required for ongoing maintenance and filter creation. Decisions regarding network visibility and monitoring devices, like network packet brokers, should be deeply researched and tested to assure that you are making a valuable investment with good longevity.

Answer another important question, can the device be operated effectively by most personnel without training and retraining? Also, can new personnel use the tool without oversight and costly time-off for training? To this end, “usability” is the key factor that allows organizations to use network equipment with ease; and still be assured that they are getting a true, reliable, and repeatable view of their traffic and network operations.

CHAPTER 2: The Impact of CLI and GUI Interfaces

The most important question is, how do you go about lowering the short term and long-term operating costs for your monitoring solution? Eliminating as much complexity as possible is the answer. There is a relationship/impact between programming complexity and network costs.

We have learned, thanks to Microsoft Windows and other graphical user interface-based systems, that a real GUI is the best and easiest way to get true repeatability in the setup and operation of network components, thus a repeatable visibility platform. To this end, a true GUI (not a CLI or CLI translator) is the best way to assure the most cost-efficient usage of any network visibility component. According to Douglas Engelbart at ARS Technica, a graphical user interface is important because it allows higher productivity, while facilitating a lower cognitive load.[1]

GUI’s have additional value in that they are hard encoded and do not slow down the network packet brokers due to the complexity and extensibility of multi-level filters or filter items, like duplicate packet deletion. Every component that requires a CLI has repeatability and program issues. For devices like switches and routers that are fully supported with command line scripts, CLI is acceptable, but it is still too complex for most IT personnel and requires training and then retraining. In fact, according to the Cisco CCNP Security Firewall 642-617 Quick Reference manual, “CLI is fast, after you have mastered it, but the GUI is intuitive and easier to configure, especially with the wizard quick-configuration options now available.”[2] Many CLI programmed visibility devices can actually can drop up to 80% of the packets.[3] This creates many false positive indications for security and troubleshooting activities, rendering the device and its output unreliable.

---

1 https://www.reference.com/technology/gui-important-95d42a64e0c41332
2 http://www.ciscopress.com/articles/article.asp?p=1681062
3 https://www.ixiacom.com/resources/tolly-network-packet-broker-test-report

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

From a human perspective, CLI can also be a significant and time-consuming source of frustration. Lack of support by upper management along with not having the needed and usable equipment have been cited as the top reasons for resignations. Losing network employees is a major security risk, and a costly event, requiring additional financial investment and loss of time to re-secure the network by changing all passwords and permissions.

Configuring devices from the command line is the time-honored tradition for network engineers. But for everyday operational tasks, the CLI is no longer fit for the purpose. As the number of devices in our networks grow, the use of the command line for operations becomes increasingly inefficient.[4] More and more companies in the network and security management world are using GUI-only configurations.

Even Cisco is moving towards GUI interfaces, which are also included in the Cisco CCNP certification![5] Examples of GUI device management products for Cisco routers and switches are: Cisco Router and Security Device Manager (SDM), Cisco Configuration Professional, Cisco Configuration Assistant, and the Cisco Network Assistant.

GUI-driven screens and drag-and-drop commands make it easy to define and launch new services. According to a Forrester Research report[6], more companies today offer low- or no-code (GUI) platforms that allow nearly anyone to painlessly interface to the technology.

Today, usability is a must-have for optimal return on investment with new technologies. Companies that focus on user-experience (UX) and user-interface designs within product and application development create better solutions, increase revenue, perpetuate customer loyalty, and improve market share. Numerous industry studies have stated that every dollar spent on UX brings in between $2 and $100 dollars in return. Already, household names such as Samsung, Charles Schwab, Motorola, Logitech, and Dell are leveraging UX and interface design in the development of their products and applications—with strong results.[7]

The features and capabilities of a programming system are the second area that will definitely affect system complexity. Let's compare CLI to GUI to see the differences. First, we need to separate the 3 types of programming that the majority of packet brokers use:

1. Command Line Interface
2. The CLI Translator/pseudo GUI
3. True GUI

---

[4] http://etherealmind.com/the-command-line-is-dying/
[5] https://www.certificationkits.com/cisco-certification/cisco-ccnp-tshoot-642-832-exam-study-center/cisco-ccnp-tshoot-complex-network-maintenance/
[6] https://smallake.kr/wp-content/uploads/2020/01/idSisdoc_12055179v2-86-Forrester-ES117623.pdf
[7] https://www.fastcodesign.com/1669283/dollars-and-sense-the-business-case-for-investing-in-ui-design

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

A command line interface is a programming method that started in the 1960's during the Teletype era. On today's computers we would use CLI by using the run command and typing in commands for the computer to carry out. How many computer owners actually do this? Very few. Again, thanks to Microsoft and other operating systems, we now use GUI's to program our systems. There are many types of CLI languages along with thousands of commands and variables. Consider that CLI programming is processor heavy and when NPB's, like Gigamon, are tested for throughput under load. Even short stage filters cannot pass 100% of the data. In contrast, the processor in a GUI environment is free to handle the filtering, collection, and passing of requested data. It often handles the display of required statistics but is not burdened by other interrupt service routines. This high level of access allows the NPB to handle a much higher bandwidth of data flow with better time stamp accuracy and deeper filtering queues. Using CLI to program SPAN sessions is difficult. This is why NetFort created a free SPAN translator called the NetFort SPAN Port Configurator that was available on their website until they were purchased by Rapid7 in 2019.[8]

The most common CLI is Cisco System's IOS. However, in a recent article, Dave West from Cisco Systems predicts that CLI will become the interface of last resort.[9] If you are still curious why, I suggest you review the basics (and complexity) that are involved in CLI programming. You can get more information from this resource.[10]

A second option is the CLI Translator, also called a pseudo GUI. This is an unusual method where the user types in, or clicks, on a desired command and has to add in the variables so that the system translates the GUI command into CLI context. This is not the best or reliable way to program. It is often preferable to use a real CLI or a real GUI over this method. There are thousands of CLI commands and variables, but this method can only handle a small and limited variety of commands and variables. That being said, there are several CLI translators and methods for this function.[11]

The third, and usually best, option is a real GUI. A real GUI is where the commands that one uses (clicks), is hard coded into the machines operating system via hard ASIC routines. Once any variables are indicated, the system is already in action. GUI's are much faster and the most repeatable methods for programming technical network equipment. In the early years, CLI was considered the most versatile interface. However, today's GUI's are actually more flexible and more repeatable and can be learned more simply and quickly. GUI programing does not require any retraining and can be used by the newest network employee.

Research[12] from the analyst group EMA shows that for the average enterprise, 74% of the respondents move or change their tool connections 2 or more times per month. For 30% of the respondents, they change their tool connections 5 or more times per month. I would assume that each one will typically have some sort of programming modifications. A second question clarified that for 27% of enterprises, IT engineers spend ¼ of their time configuring monitoring tools. Another 28% spend up to 50% of their time configuring tools. And another 20% spend up to 75% of their time configuring tools. What this means is that the time it takes to program, or reprogram a monitoring filter, will directly affect your total cost of ownership. The programming of data filters isn't a one time, or once a year, activity. It's an ongoing activity you will want to account for ahead of time when performing a TCO analysis.

---

[8] NetFort SPAN Port Configurator - YouTube
[9] http://www.theregister.co.uk/2016/03/10/cisco_says_cli_becoming_interface_of_last_resort/
[10] https://en.wikipedia.org/wiki/List_of_CLI_languages
[11] https://en.wikipedia.org/wiki/List_of_command-line_interpreters
[12] An EMA research survey commissioned by Ixia in October 2016

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

Years ago, I wrote an article for www.lovemytool.com comparing CLI to the Keysight (Anue at that time) GUI. In the article, I wrote a simple CLI script to deliver the traffic to and from a certain IP address. This article has been read over 2,000 times by various competent network engineers and, as of yet, no one has reported to me that the script I wrote would not work, even though I purposely included some errors. This shows how easy it is to overlook a small mistake that renders a simple, but still multi-line, script useless.

CHAPTER 3: An Analysis of Tap and SPAN Technology

The type of network monitoring equipment deployed also affects complexity. IT managers should be especially interested in two types of equipment—data access and data filtering. Regarding the first category, the two most common ways of accessing monitoring data are through either a switched port analyzer (SPAN) port or a Tap. Let's examine both methods.

A common way of capturing network data for monitoring purposes involves the use of SPAN ports, also called mirroring ports. These ports are typically available off of a network routing switch. A SPAN port should not be confused with a SPAN session. A SPAN session is a CLI monitor command, or set of commands, used to create a basic filter. The SPAN port is still the main monitor access mechanism for the switch bus.

While SPAN ports make a mirrored copy of network data, there are a host of issues associated with them. This needs to be factored into your monitoring strategy. For instance, the use of SPAN ports creates the following issues:

- Duplicate data packets are created which reduces the efficiency of your monitoring tools
- Missing data (Layer 1 data, corrupted and malformed packets, bad CRC, Interframe gap, and other data oddities) is not forwarded on to SPAN ports. Therefore, SPAN access is not suitable for real time protocol (RTP) monitoring, capture, and analysis, especially in modern mean opinion score (MOS) and quality of experience (QoE) strategies.
- SPAN ports only provide summarized data
- SPAN ports change the time stamps of packets
- SPAN ports have been shown to be hackable (so they can be a security risk)
- SPAN ports require CLI programming

In fact, SPAN ports themselves are one of the reasons you can develop network blind spots. Depending upon how you set up the filtering (i.e. what traffic you decide to make a copy of and route to the SPAN port), you may be collecting the wrong data and/or accidentally clipping (i.e. dropping) data you are actually interested in. To sum it up, you're not seeing a complete copy of the traffic on your network.

One question that comes up is whether SPAN/MON port access is a passive technology? The answer is a resounding No! While some people try to call SPAN port technology a passive data access solution, passive means "having no effect" and spanning port (mirroring) does have measurable effect on the data.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

Here are some of the ways that a SPAN session modifies monitoring data:

1. Spanning changes the timing of the frame interaction (what you see is not what really happened)
2. The spanning algorithm is not designed to be the primary focus (main function) of a network switch, like switching or routing is. So, the first priority is not spanning and if replicating a frame becomes an issue, the hardware will temporally stop the SPAN process.
3. If the speed of the SPAN port becomes over loaded, frames are dropped
4. A SPAN port drops all packets that are corrupt as well as those that are below the minimum size. So, all of the frames are not passed on. No Interframe data is passed, either. All of these events can occur without any notification being sent to the user. This means there is no guarantee that one will get all the data required for proper analysis.

Proper spanning, even if the port could handle the load, requires that a network engineer configure the switches. This takes away from the more important tasks that network engineers have. In addition, many network configuration changes can become a political issue due to creating contention between the IT teams, the security teams and the compliance teams, etc. SPAN programming can also require Change Board approval, which introduces data capture delays.

When we only had 10 Mbps links and a robust switch (like ones from Cisco), engineers could almost guarantee that they could see almost every packet going through the switch. With 10 Mbps fully loaded at around 50% to 60% of the maximum bandwidth, the switch backplane could easily replicate most of the frames. Even with 100 Mbps, one could be somewhat successful at acquiring most of the frames for analysis and monitoring. And if a frame or two here and there was lost, it was no big problem.

This has all changed with 1 Gigabit to 100 Gigabit technologies; starting with the fact that the maximum bandwidth is now twice the base bandwidth. A full duplex (FDX) Gigabit link is now 2 Gigabits of data and a 100 Gigabit FDX link is now 200 Gigabits of potential data. No switch or router can handle replicating/mirroring this amount of data for all its ports, plus handling its primary job of switching and or routing. It is impossible to pass all frames (good and bad, including FDX traffic) with any real-time correlation from a SPAN port.

A Cisco white paper on SPAN port usability and the use of SPAN port for LAN analysis[13], warns that “the switch treats SPAN data with a lower priority than regular port-to-port data.” In other words, if any resource under load must choose between passing normal traffic and SPAN data, the SPAN port loses out and the mirrored frames are arbitrarily discarded. This rule applies to preserving network traffic in any situation. For instance, when transporting Remote SPAN (RSPAN) traffic through an Inter Switch Link (ISL), which shares the ISL bandwidth with regular network traffic, the network traffic takes priority. If there is not enough capacity for the remote SPAN traffic, the switch drops it. Knowing that the SPAN port arbitrarily drops traffic under specific load conditions, what strategy should users adopt so as not to miss frames? According to the Cisco paper, “the best strategy is to make decisions based on the traffic levels of the configuration and when in doubt to use the SPAN port only for relatively low-throughput situations.”

[13] http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/san-consolidation-solution/net_implementation_white_paper0900aecd802cbe92.html

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

It should be noted that there are times when spanning is okay. For instance, many monitoring events can, and do, successfully use spanning as the packet access technology. These monitoring events are looking for low bandwidth application layer (ULP) events like “conversation or connection analysis”, “application flows”, and applications where real-time (and knowing real delta times like voice and video flows) is not an important factor or requirement. SPAN ports can be used for the inventory of addresses and other non-time sensitive monitoring, which today is a very limited view of our complex applications and network traffic. However, SPAN ports are NOT acceptable for today’s security monitoring applications and modalities.

The monitoring requirements just mentioned utilize a small amount of bandwidth and packet grooming. This means these packet drops do not affect the quality of the reports and statistics. The reason for their success is that they keep within the parameters and capability of the SPAN ports ability. These specific applications do not need every frame for their successful reporting or analysis. In other words, if used correctly, SPAN ports are a usable technology as part of a well-managed methodology.

In summary, the fact that a SPAN port is not a passive data access technology or even entirely non-intrusive can be a problem, particularly for data security compliance monitoring or lawful intercept. Since there is no guarantee of absolute fidelity either in time or actual packets, it is possible, and even likely, that evidence gathered by this monitoring process will be challenged in the court of law.

The other access technology is called tapping. Taps ARE passive devices that can be installed anywhere in the network to give you access to all of the data at that location. This is different from a SPAN port that can only give you access to what is available from the core switch. Taps can be installed in the core, at the edge of the network, or anywhere there is a need (like some type of perceived bottleneck). When combined with an NPB, you can pool monitoring resources to maximize the efficiency of existing tools (through aggregation and load balancing across multiple tools).

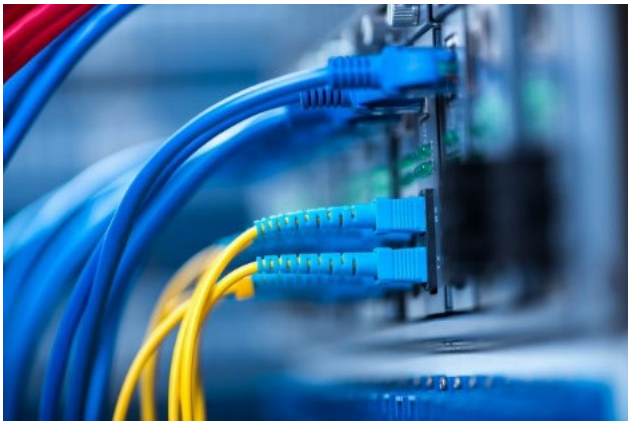

Taps offer significant advantages over the use of SPAN ports to monitor the network. SPAN ports require an engineer to configure the ports on a network switch. Taps do not. A Tap also passes all data on a link, including the Interframe gap, errored packets, and short and long packets—all with a REAL time stamp and in the order of arrival. This includes bad frames that can be caused by a faulty network interface card (NIC) and duplicate packets.

Another benefit of Taps, since they are passive, is that they do not affect frame timing. Any active device that touches a frame has changed the frame timing—even if nothing more than changing its absolute timing reference to the network. It is essential to keep all changes by a device linear. If the frame offset was 10 ms, then all frames should have the same offset. If not, the device is interfering with the real-time analysis capability of that access point. SPAN access is a great example of variable offset and the impossibility of doing authentic time-based analysis from a SPAN/monitor port. A good Tap with a tested algorithm handles the Send and Receive integration with consistent timing for the best visualization. All access devices can change the frame and its environment. However, as long as the company providing it and the operator understands this, then one can get relevant data and facts from the devices.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

A Tap is the ONLY device that will pass every bit, byte, nibble and octet. This includes the Inter-frame gap, bad, large, small, and other error packets which are needed to properly monitor and troubleshoot all network problems. Even if one uses a higher technology filtering device, it is strongly suggested that you stick with using a Tap as your media access. This means a standalone Tap, not an integrated one. It should be noted that there is significant debate about the viability of passing bad packets for capture and post capture analysis. I feel that just counting the bad packets/types IS acceptable, and in fact, a requirement for baselining analysis purposes.

A detailed comparison of Taps and SPAN ports is as follows:

- Taps do not alter the time relationships of frames – spacing and response times especially important with VoIP and Triple Play analysis including FDX analysis.
- Taps do not introduce any additional jitter or distortion which is important in VoIP / Video analysis.
- VLAN tags are not normally passed through the SPAN port, so this can lead to false issues detected and difficulty in finding VLAN issues.
- Taps do not groom data nor filter out physical layer errored packets.
- Short or large frames are not filtered.
- Bad CRC frames are not filtered.
- Taps do not drop packets regardless of the bandwidth.
- Taps are not addressable network devices and therefore cannot be hacked.
- Taps have no setups or command line issues, so getting all the data is assured and saves users time.
- Taps are completely passive and do not cause any distortion, even on FDX and full bandwidth networks. They are also fault tolerant.
- Taps do not care if the traffic is IPv4 or IPv6. They pass all traffic.

There is one more major and important consideration about access technology. Do not forget that any access device can be called into question in civil and criminal cases. When using the data captured as evidence in employee misuse or for CALEA/lawful capture type situations, a Tap is your very best ally. It presents the evidence with NO CHANCE of changing anything and it provides a solid timing reference. This is called forensically sound data/evidence and is mandatory for court-ordered evidence. Another advantage to consider in our security conscious world is that a real Tap cannot be hacked, as it does not have an IP address. A Tap is the only truly secure way to access and monitor your network. Therefore, any evidence gathered with the device is as pure as it can get.

More information on Taps versus SPAN/MON and VACL's is available at this website.[14]

[14] https://www.networkdatapedia.com/post/2018/01/01/the-original-tap-versus-span-comparison

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

CHAPTER 4: Outdated Processes Are Costing You Money

As described in the previous sections, there are two common practices for network monitoring that are outdated. The first is the use SPAN ports and the second is the use of a command line interface. The previous sections detailed these concepts, including benefits and detractions. But what are the costs associated with the use of these two practices? We'll discuss those consequences here.

An initial reason often cited for using SPAN ports is that they are "free". Most network switches have two SPAN ports included. However, using SPAN ports is actually costing you money. This comes about because of the inaccuracies that SPAN ports introduce into the monitoring network. As mentioned previously, SPAN ports do not deliver an exact copy of the data and data flows. They also deliver only parts of the data and might not deliver any of the data if the network switch is heavily loaded. This results in missing data, duplicate data, and real problems that have become obscured and hidden. Please remember that a switch was designed to handle and deliver packets to the correct group, thus offloading thousands of packets to be read by NIC's unnecessarily. A switch is analogous to the post person delivering your mail to the correct access without one having to sort out what is your mail and what mail belongs to someone else on another street. Switches were never designed to be network monitors so monitoring functionality has the lowest priority. SPAN and MON access came from an old quality control test that required access to the data bus.

In addition, SPAN ports are often the source of duplicate packets, thus adding more issues to troubleshoot. For instance, this can be an issue if one believes that the switch is causing the duplicate packets and decides not to investigate any further. Later on, that engineer discovers that there are really duplicate packets in the network. The existence of that duplicate data can indicate there is a major issue, like a failing piece of equipment, an architecture issue, or potential malware.

The first two results (missing and duplicate data) force your security and monitoring tools to work harder and can even make them less effective. More CPU cycles are spent on irrelevant tasks, especially in the case of removing duplicate data. For security tools like an IPS, "session stickiness" is often required. Missing data can have the result that the security tool does not detect that a session has closed. If too many sessions remain open, the tool's memory can't track any more sessions. In some cases, the security tool will shift from an "inline blocking mode" state to an "out-of-band detection mode" state. It then sends a trouble alert but ignores additional sessions, allowing them to pass downstream without inspection. This means that the device isn't actively analyzing those potential security threats. It can also be a manual process for an IT engineer to issue a command to move the tool back into an inline state. Then the engineer needs to perform some sort of analysis to see what triggered the incident, which costs more time and effort.

Other possible effects from missing data can include false positives for threat detection and troubleshooting activities, longer time to resolution intervals which results in longer mean time to repair (MTTR) objectives, and a longer amount of time for security and monitoring tools to analyze the data. False positives for troubleshooting solutions cost additional time and money and leads to a lack of confidence in the monitoring solution. SPAN data is never in relative time to real network events. This is a diagnostic killer for all RTP events like voice, video, and any other time-based measurement.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

The common practice of using a CLI also has higher costs. For instance, the time to create a data filter within an NPB can be four to ten times faster than using a CLI. Data from a study I conducted shows that CLI programming can result in an additional $6K annual cost over simply installing a Tap and NPB combination.

The comparison below is an attempt to perform an “apples to apples” comparison with regards to SPAN ports and Tap ports. While there are several technical reasons that Taps are superior to SPAN ports, there is almost always a financial discussion that takes place where the customer states that SPAN ports are free and that while Taps may have some technical benefits, the costs don’t outweigh the benefits. While the cost of a Tap is only about $629 street price, the financial analysis provided below can be used to further explain the true costs of using SPANs.

Customers typically already have Cisco switches with 2 free SPAN ports included. However, proper mirroring requires a network engineer to configure the switches properly. Configurations can require frequent changes, depending upon priorities and frequency of monitoring (IT team data needs, the Security team needs and the compliance team needs), as the number of available SPAN sessions is often limited. If SPAN sessions are limited, and the IT engineer insists on using SPAN ports, then additional SPAN ports will need to be added by buying more Layer 2 switches or converting switching ports to mirror ports.

The hardware analysis below shows the cost of adding a “representative” SPAN port vs. adding a “representative” Tap port. At the very low end, a SPAN port can be more attractive financially than a Tap port, based upon the hardware cost. In the mid and high ranges, the hardware cost advantage wains and finally goes away at the high end. This comparison uses a Keysight Flex Tap vs several Cisco switch models. The Cisco switches chosen are very subjective and the SPAN cost per port could actually be much more expensive than what is shown below.

Hardware Cost Comparison

Cost Model	Tap Cost (per Port)	SPAN Cost (per port)
Low end (10G)	$629	$273
Medium (10G)	$629	$477
High end (40G)	$629	$1,111

The real costs of SPAN ports come from the management overhead. This is shown in the following chart. The cost to administer a Tap is typically $0. It is a one time “set and forget” process that needs no configuration and may not even need external power, which is the case for the Keysight Flex Tap. In contrast, administration costs for SPAN sessions start Day 1, as illustrated below. **SPAN administration costs are a hidden cost of doing business that is often not appreciated by IT personnel.** In the conservative example shown below, the average annual recurring maintenance costs ($6,890) for SPAN sessions could have been redeployed to buy an average of 10 Taps (annually).

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

Configuration Programming Cost Comparison (for 1st year)

Provisioning	Tap Cost	SPAN Cost
Initial Set-up	$0	$530
Session 1	$0	$97
Session 2	$0	$302
Session 3	$0	$540
Session 4	$0	$864
Session 5	$0	$957
SPAN session planning	$0	$3,600
Averaged Total	**$0**	**$6,890**

What is not shown, but should be equally concerning, is that there will be additional delays in Mean Time To Repair of network problems if you use SPAN ports. SPAN ports are part of the network. As such, any configuration changes can affect the delivery of information to security and monitoring tools, as well as data loss on the corporate network (if SPAN ports are over provisioned). For these reasons, SPAN ports are usually configured outside of working hours. However, for troubleshooting purposes, SPAN ports may need to be reconfigured during business hours (say if the network is running slow or parts of it are down altogether). This will require a Change Board approval which can typically take 5 to 6 hours to acquire. This delay cost needs to be accounted for as well.

In addition, CLI filters are prone to errors since they are often manually created. This creates a significant potential source of errors and debugging time required to troubleshoot those data filters. And while some errors are obvious upon review, others are not and may result in clipped data that delivers some of the data (but not all of the requisite data) to the security or monitoring tool. Again, this results in erroneous conclusions and delays in time to resolution.

NPB's that have built-in filter creation engines can remove this issue for IT and security managers. Once the filter is created, it needs to be validated. This can take over an hour to validate the filter. NPB's with built-in filter engines can often validate themselves. If not, a one-time validation process to prove the filter engine accuracy should be enough. Each individual filter doesn't need to be validated like it does when the filter is created through a CLI process.

I decided to perform a comparison myself between a CLI and a GUI. This analysis was performed using a command line interface to set up a simple data filter for deleting SSL encrypted data. This showed that a GUI interface was about five times faster than a CLI interface. This analysis was based upon a Cisco Catalyst 6500 switch and a Keysight Vision ONE NPB. The time to literally set up and execute a filter using the Keysight interface was about 2 minutes and about 15 minutes for the CLI interface. This data can be extrapolated to create a financial analysis of CLI vs. a GUI.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

Performing a side-by-side comparison of CLI costs versus GUI costs, I can confidently say that a GUI can cut your long-term operating costs by at least 75%. This is based upon the following assumptions:

Cost Component	CLI	GUI	Frequency
System Maintenance	$6,000	$2,000	Annual
Training/Retraining	$5,000	$0	Annual
Initial/annual Filter setup	$6,000	10 hours per year	Annual
Normal filter changes	15 mins per filter	2 mins per filter	3 times/month
Troubleshooting incident	4 hrs per filter	15 mins per filter	4 times/year
Labor rate	$100/hr	$100/hr	

This data yields the following analysis:

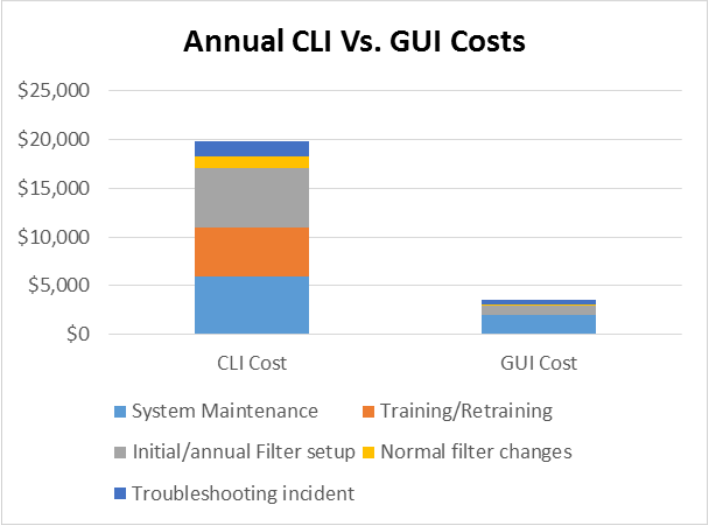

Here is a table with the detailed costs:

Cost Category	CLI Cost	GUI Cost
System Maintenance	$6,000	$2,000
Training/Retraining	$5,000	$0
Initial/annual Filter setup	$6,000	$1,000
Normal filter changes	$1,200	$120
Troubleshooting incident	$1,600	$400
Total	**$19,800**	**$3,520**

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

CHAPTER 5: Network Packet Broker Vendor Summary

Network packet brokers are great products and I highly recommend them. They perform a value-add function when inserted between the monitoring data access points (Tap or SPAN) and the monitoring tool(s). Here is a short list of the value that an NPB can add:

- Gain link-layer visibility and data access across entire network
- Centralize tools while increasing their reach and greatly improving tool ROI
- Provide flexible access to both passive and active inline tools
- Boost monitoring and security tool efficiency
- Reduce both CAPEX and OPEX through longer tool lifecycles
- Support network upgrades by load balancing existing tools
- Quickly provision new tools by eliminating SPAN port contention
- Centrally, remotely, and/or locally manage network visibility and access

There are several NPB devices for sale in the market today. No matter which NPB you choose, one should always consider these main points:

- Will the device deliver the visibility and access you need, today and tomorrow?
- The cost and time for training, and retraining
- The time required to set up and activate complex filters
- Everyone on the Network and Security Team should be able to use the NPB
- Versatile visibility should not be paid for by complexity
- No event should cause loss of packets or the original packet time stamp
- Will the company behind the NPB be around a long time?
- Does the company have a long history of great customer support?
- Is the management team known for honesty?
- How fast can you implement a filter? This is the true key to a successful deployment.

So, who are the primary players in the network packet broker space? The three largest vendors are:

- Gigamon
- VSS/Netscout
- Keysight

I reviewed published materials on the websites and industry publications for these three vendors. Based upon this material, I came to following conclusions regarding each vendor.

Gigamon Systems has been around since 2004. They make a line of packet brokers and Taps. Their packet broker has the capability of performing various NPB functions including: aggregation, Layer 2 – 4 filtering, deduplication, and basic Layer 7 filtering. Specifically, I looked at the GigaVUE-HC2 product.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

From a technical point of view, the product can perform many basic functions and a few advanced functions. However, the product has several noticeable limitations including:

- Business Wire reported that the Tolly Group tested the Gigamon NPB against another vendor (Keysight) and found that Gigamon's NPB dropped packets under various conditions and didn't report the losses. See this report for details.[15]
- Gigamon documentation shows that there is a feature combination map needed to understand which features can be turned on at the same time within a single module. This means that all features cannot run simultaneously at line rate.[16]

In addition, this vendor uses a combination of a CLI interface and a CLI Translator (Flowmap) to create filters. As mentioned earlier, this type of interface has quite a few drawbacks including complexity. For instance, some application definitions are defined to help with the Layer 7 filtering capability but any other applications need to be defined by the customer. This includes creating the filter, validating it, and performing application signature updates as the manufacturer changes their application. This requires significant maintenance activities.

A second packet broker vendor is VSS, which was acquired by NetScout Systems in 2014. The VSS vBroker product uses policy-based filtering. This is essentially a mid-market solution as it is not keeping pace with the industry needs.

In looking at this product, I had the following concerns:

- It performs packet grooming. An NPB should not groom packets.
- It does not perform true load balancing. Up to eight monitor ports can be configured to output traffic as a single logical pipe, with their output ***approximately*** evenly distributed throughout the load balanced group

NetScout does have another solution, the nGenius Packet Flow switch. This is part of the nGenius Service Assurance Solution which is focused and designed for their high-end system-based solution. This solution is typically sold as part of their network monitoring and analysis system and is not usually sold as a separate solution. The VSS vBroker is the standalone offering from NetScout.

The third vendor is Keysight. Keysight has been in business since 1997, mainly as a test company. However, Keysight acquired two packet broker solutions over the last several years. They acquired the Ixia (Anue) NTO product line and the Ixia (NetOptics) Tap and packet broker line in 2017.

I looked at the Vision ONE packet broker solution, which is based on the Ixia (Anue) NTO product line. This product has all of the features that I would consider basic NPB features (aggregation, load balancing, Layer 2 through 4 filtering, and regeneration). It also has advanced NPB features including deduplication, packet slicing, time stamping, header stripping, Layer seven filtering, NetFlow generation, geolocation of users, and several more. All features are supported concurrently and at line rate, up to 40 Gbps in this model.

The product uses a GUI interface for everything. No CLI is required, or offered. The GUI is quick, simple, and intuitive. No filter rules were needed. It simply uses a drag and drop, point and click technology to quickly create filters and activate them. Activation of the filters were

[15] http://www.businesswire.com/news/home/20160120005497/en/Testing-Demonstrates-Ixia-Network-Packet-Broker-Delivers

[16] https://docs.gigamon.com/pdfs/Content/Resources/PDF%20Library/GV-51600-Doc/GigaVUE-FabricManagement-Guide-v5.16.pdf, p. 1296-1299

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

completely simple and hitless. The product also helps the user complete even the most complex filter including Boolean layers and the units test the configuration before completion. This is a huge step over any other state-based setup process. When it comes to the layer 7 filtering, Keysight has almost 250 application signatures defined. This is an extensive amount and significantly diminishes the need for the customer to create their own application signatures.

The Business Wire article mentioned earlier reported that the Tolly Group tested the Keysight NPB against a CPU-based NPB and that in all its tests, the Keysight NPB demonstrated that the packet processing performance of the Vision ONE (NTO) product line does not change based on configuration. Keysight delivered 100% visibility across all network operating conditions tested. A copy of the test report is available on the Keysight website.[17]

To spend lots of time comparing different vendors would be a waste of time as everyone focuses on their special area, mostly made of mundane comparison points. I have used almost every type of NPB, and I like the simplicity of a GUI interface and the repeatability and reliability for filters that I create. Even if I have a stored CLI filter to use for different variables, I still can make simple mistakes that cause filter failure.

I have heard from several buyers of NPB's that say that they bought A, B or C vendor's tool, mostly based on some level of concern that they needed the associated complexity to assure their management team that they would be able to deploy some magic filtering capability. Most of these people have gone back and purchased the Keysight NPB saying that if they needed some future support, they had it but also, they could use this NPB so much easier and with confidence. They could teach anyone how to use it in a very short time and the GUI gave even the newest Team member easy access to build filters for network, application and security visibility studies. After all, that is what one needs in a NPB—ease of use developing filters and a deep understanding what they are getting on their different focused tools or deep packet capture engine.

To me, that is all the comparison one needs. I liken the Keysight NPB to a calculator and the other two NPBs to slide rule devices. The Keysight Vision ONE (NTO) was the first bottom-up designed filter engine using a hard platform, not just a switch SPAN chip running on a basic processor bus. Years ago when I was investigating packet brokers, I came up with this mantra – "Simplicity is the Goal of Advance Technology". I believe that Keysight's real GUI really fits this mantra.

As mentioned, I tested the Keysight NTO product. In early 2016, Keysight introduced its newest packet broker product, the Vision ONE. This NPB is based upon the original NTO product set and combines inline and out-of-band monitoring capabilities. Keysight uses a patented Dynamic Filter Engine to make data filter creation quick, simple, and hitless (i.e. no packets dropped). This Dynamic Filter Engine automatically enables multiple, overlapping rules to be applied transparently. The net impact of the Keysight technology is that time spent on filter planning, creation, and maintenance is dramatically reduced.

The Vision ONE user interface is simple to understand and easy to deploy. It's based upon drag and drop and point and click technology—no CLI needed and no CLI expertise needed. You just create filters for what you want to see. Pre-configured, hard coded filters (also called floating filters) and easy to follow filter libraries in the Keysight solution allow network personnel to create multi-layer and complex Boolean diagnostic filters, attach them to tools, and place those filters in standby/saved mode for when you need them for fast and easy deployment. When a network issue or event arises, it takes less than a minute to attach a

[17] https://www.ixiacom.com/resources/tolly-network-packet-broker-test-report

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

floating filter to a network port and to a specialized tool to begin diagnostic capture and analysis.

Which of these programming scenarios would you prefer?

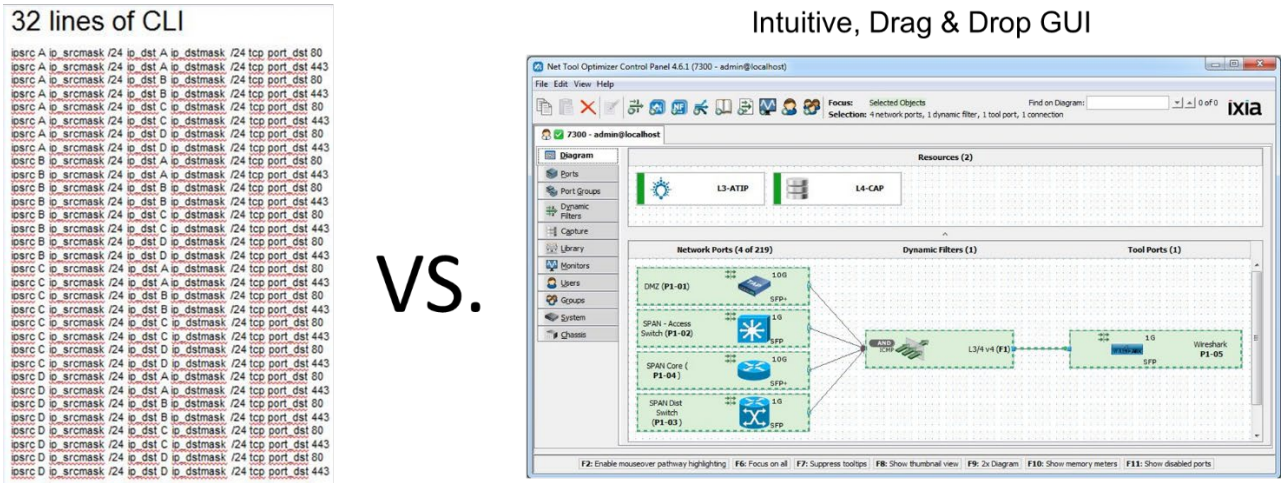

Mike Pennacchi of NPS (http://www.nps-llc.com/), a well-known and highly respected Network Developer, Data and Incident Analyst and Network technology instructor, stated in a conversation with me (a very long time friend and associate), “CLI programming is at least 5 times more difficult than GUI programming in the packet broker filtering arena.” Mike said that, even with saved CLI programs, when one is making a correction or addition for another usage, even one wrong character (a space dash, underline, etc.) can cause the filter not to work and thus not acquire the needed information.

A head-to-head comparison that I ran shows that creating filters using the Keysight Vision ONE interface is 4+ times faster than Gigamon. This saves you time so you can work on other tasks. Keysight’s patented Dynamic Filter Engine removes potential data filter misconfiguration and data clipping errors. This saves time, money, and speeds up troubleshooting time to resolution. As most IT managers will attest, troubleshooting time to resolution is extremely important because it is directly associated with measurable internal and external business costs. This is backed up various industry studies including the following:

- According to a recently commissioned survey by Talari Networks, 89% of IT professionals have had at least one unplanned outage this year with 69% having encountered two or more unplanned outages in the last twelve months.[18]
- A report from IHS found that midsize to large companies typically experience five minutes of downtime every month, which translates to a cost of about $1 million annually for midsize firms and $60 million a year for large enterprises.[19]
- Infonetics found in its ICT downtime survey that businesses lose nearly $4 million a year to network downtime, which equates to 0.5 per cent of a company's total revenue.[20]

[18] http://www.itproportal.com/features/what-does-a-network-outage-really-cost/
[19] https://www.businesswire.com/news/home/20160125005188/en/Businesses-Losing-700-Billion-a-Year-to-IT-Downtime-Says-IHS
[20] https://www.cablinginstall.com/data-center/article/16472303/report-ict-downtime-costs-businesses-4-million-per-year

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

- The Ponemon Cost of Data Center Outages Report form January of 2016 found that the average Data Center down time cost is approximately $7,793 per minute! This uses their figure of an average outage cost at $740,357 per incident with an average outage duration of 95 minutes per incident.[21]

If you can save even 10 hours a month and have a quicker response to network events, this is an average savings of almost $4,675,800. Plus, every time you fail to easily find and respond to issues, one's professional standing can be harmed. No one ever wants to be known as the person that took a month to find a data leak because they were unable to get the right data to the correct tool because of programming issues with their NPB. Getting the data should be the easiest task and analyzing the data should be the major professional task. Remember – downtime and network issues are also an emotional issue for all involved in finding and mitigating those issues. This is an unmeasurable human cost!

The drag and drop interface of the Keysight product is intuitive to use and understand. No training courses are needed. This means your system is up and running in a minimal amount of time. Keysight has at least one documented case study where one of their customers was up and running in less than 30 minutes. Other vendors take much longer. Again, this was one of the initial criteria I laid out as being important to the overall TCO.

CHAPTER 6: Conclusion

IT managers are looking for monitoring solutions that provide the following:

- Greater return on IT investment, by dramatically improving the effectiveness of security and monitoring tools
- A solution that allows the Network Management and Security Teams network visibility access allowing them to be successful in quick recognition of major events and allowing them to use focused tools for rapid mitigation of events
- Enhanced security by eliminating network blind spots that could be concealing intrusion attempts, signs of abnormal and or bot traffic, or data exfiltration following a successful exploit
- An in-depth understanding of traffic volumes with complete data monitoring to help predict when systems may be about to fail enabling IT and security teams to gain control of their networks

The best solution I have found for data filtering is the Keysight Vision Series product line. Keysight's visibility solution, Vision ONE, has been proven to be easy to install, easy to configure, and easy to maintain. The ease of use of this solution lowers the TCO for any visibility solution purchase. This network packet broker is by far the easiest to use plus has the most advanced filtering capability with tested and proven repeatability, shown to be the fastest in the industry by the Tolly group. The Keysight products meet all full duplex bandwidth requirements up to 100 Gbps with the highest density of ports with full bandwidth even with deep and complex filters.

[21] http://files.server-rack-online.com/2016-Cost-of-Data-Center-Outages.pdf

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

Keysight’s patented Dynamic Filter Engine accomplishes this by making filter creation quick, simple, and hitless (no packets dropped). The Dynamic Filter Engine automatically enables multiple, overlapping rules to be applied transparently and eliminates issues with filter clipping and misconfigurations that often result from command line interface attempts. This guarantees the accuracy and repeatability of your filters, whether you are filtering two ports or two hundred. This reduces engineering time spent on filter planning, creation, and maintenance.

Why spend your precious time debugging CLI-based filters? All of the complexity of designing a filter, creating it, and testing its accuracy has been removed from the user’s cares.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS

A Note About The Author

Tim O'Neill is an independent network and communications technology consultant. Tim started his technology career when he became a Ham Radio operator at age 13 and has never quite been able to leave the world of Electrical and Electronic communications technology. Tim has spent 45+ years working in the RF, WAN, Analog, ISDN, ATM and LAN/Network analysis and monitoring market.

He has been an executive manager and a technology director resulting in many successful products and strategies for companies like Shell Research, Litton Industries, Spectron, Navtel, Network General, Ganymede and ClearSight Networks. Tim has been a technical consultant on several box office movies and has been in the high-tech industry at all levels from basic engineer to senior executive. He helped bring to market the first Digital DataScope with some of the very first digital TAPs (RS-232) in 1976 and has since designed or collaborated on many subsequent technology advancing products and test philosophies.

Tim is a Senior member of the IEEE, InfraGard, Digital Forensics Association, Computer Security Institute, ISSA, SEG, Intertel, IACSP, and other Technology and Professional societies and has served on the ANSI T1, M1 and E1, IPNG and ATM standards committees. He is an advisor and regular presenter at Kennesaw University, The Center for Information Security Education. He is also a Georgia POST (Police Officer Standards and Training) certified Instructor in the Cyber Security and Data Forensics and is a Georgia Bureau of Investigation Internet Safety instructor. He carries many certifications from WiFi to network and computer forensics, and e-discover. He was the 2014 keynote speaker at Sharkfest, and has spoken and trained for ISSA, InfraGard, Interop, MiniMicro computing and several other technical shows and conferences.

Tim serves on corporate boards and as technical advisor for several high-tech companies. He is also the Chief of Technology and Contributing Editor for www.networkdatapedia.com, a very successful website designed to help network managers/engineers gain access to valuable network technology information with real solution stories from customers, real users and technology companies with REAL World solutions.

COPYRIGHT © 2017-2022 B.T. SOLUTIONS