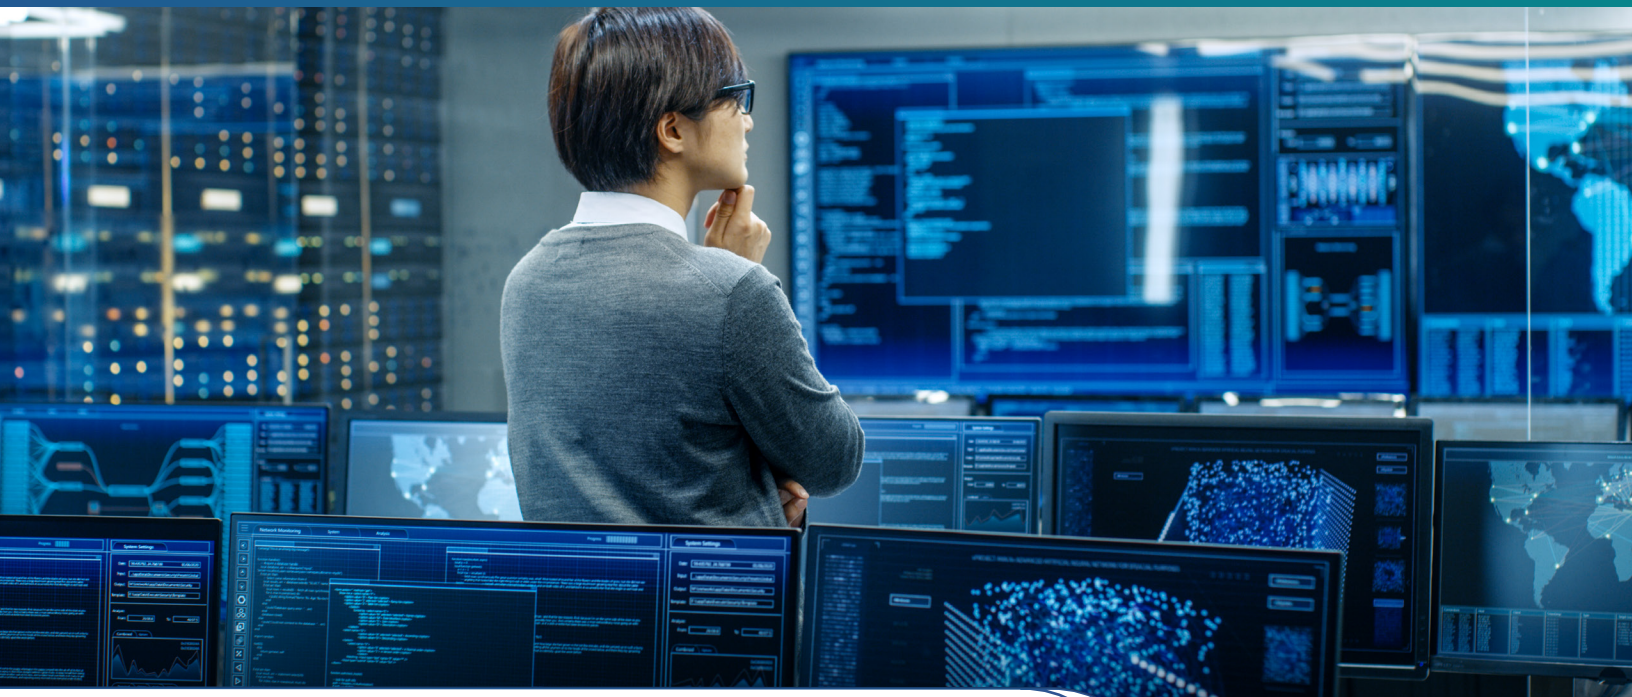

WHITE PAPER

Deploying a Layered Visibility and Cybersecurity Architecture

Fundamental Shift in Government Cybersecurity Occurring

It should be no surprise to anyone that government data networks are under a constant threat of attack. According to a Wired Magazine article (*Inside The Cyberattack That Shocked The US Government* by Brendan I. Koerner), the United States Office of Personnel Management (OPM) website alone is attacked approximately 10 million times per month. However, the discovery of the April 2015 OPM security breach that compromised the personnel records of 80 million Americans was a wakeup call that something needs to be done.

President Trump's Presidential Executive Order 13800 mandates all government agencies to adopt the National Institute of Standards and Technology (NIST) *Framework for Improving Critical Infrastructure Cybersecurity*. Part of the new executive order mandates government networks to operate within a context that supports an Identify, Protect, Detect, Respond, and Recover approach. This includes all of the network, not just security-focused appliances and solutions.

As part of this initiative, there are three actions to consider for strengthening network security and eliminating component obsolescence of critical infrastructure components:

- Understand why your security and visibility architectures should be integrated
- Understand the concept of security resilience and how it helps cybersecurity
- Deploy a visibility architecture to support the NIST cybersecurity initiative

President Trump's Presidential Executive Order 13800 mandates all government agencies to adopt the NIST Framework for *Improving Critical Infrastructure Cybersecurity*.

KEYSIGHT TECHNOLOGIES

Visibility is Security

One of the biggest security challenges for information technology (IT) staff today is to get the proper network information they need, when they need it, so that they can make informed decisions about network security and problem resolution. Proper network visibility is the solution. Without this visibility, how can you be certain your network has not been breached? If your network has been breached, what was affected? This lack of visibility gets back to the fundamental concerns that government chief information officers (CIOs), chief information security officers (CISOs), and information systems security officers (ISSOs) have. According to Rob Lewis, Enterprise Threat Manager for the U.S. Department of the Interior, "If we can't see the adversaries and what they are doing on the network, then it's very hard, almost impossible to make an impact and actually secure the data and the networks for the Federal government."

"If we can't see the adversaries and what they are doing on the network, then it's very hard, almost impossible, to make an impact and actually secure the data and the networks for the Federal government."

— Rob Lewis

Enterprise Threat Manager for U.S. Department of the Interior

The following fundamental mind shift must happen for government CISOs, ISSOs and their teams: *network visibility is network security*. Security tools and technologies are only as good as the network data they receive for analysis. If you cannot recognize the threat, then you cannot defend against it. If you cannot recognize the intrusion or breach, then you cannot stop it and mitigate the damage. And if you don't know what systems were affected, you cannot be certain you've recovered to a secure state. A best practice is to also integrate your security architecture with a network visibility (monitoring) architecture. Organizations often treat these areas as silos, which starts a cascade of problems, like process failures, blind spots, missed critical data, and delays in problem resolution.

For example, the obvious consequences of this lack of visibility are security breaches. If your agency is anything like your civilian counterparts, then there is a better than 50 percent chance that you will also not find the breach yourself. This means that your agency is informed by citizens (customers), other agencies (like the Federal Bureau of Investigations or the Department of Justice), or civilian contractors and businesses that a breach has occurred. This will be a huge embarrassment both personally and for your government agency.

Other consequences include a high number of security alerts and false positives, alert fatigue for your staff, missed alerts, and higher costs and timeframes to mitigate the damage from attacks and breaches. There is no perfect security product that can stop intruders in their tracks. It needs to be a concerted process of best practices that are put into place and maintained.

The first step is to put a visibility architecture in place that supports your security plan. A visibility architecture is essentially a cost-effective design that provides access to network traffic, intelligently filters the requisite data, sends the groomed data to analysis tools, and then delivers information as output from the monitoring tools so that IT can make informed decisions about problem resolution and network improvements. With the proper visibility architecture in place, you'll be able to see what is (and what is not) happening on your network.

The first step is to put a visibility architecture in place that supports your security plan. A visibility architecture is essentially a cost-effective design that provides access to network traffic.

Once a joint security and visibility architecture is in place, it will provide three valuable attributes to mitigate your security threats:

- Better data to analyze security threats
- Better operational response capabilities to attacks
- The ability to apply a set of consistent policies across your network

The basis of a visibility architecture is simple. First, you need to deploy test access points (taps) to get data where and when you need it. Second, you need to deploy a network packet broker (NPB) to perform data aggregation, regeneration, filtering, deduplication, and other functions to refine the data subset to exactly what each type of security tool needs. The last step is to deploy whatever mixture of purpose-built security and monitoring tools that you need and supply those tools with the filtered data from the NPB.

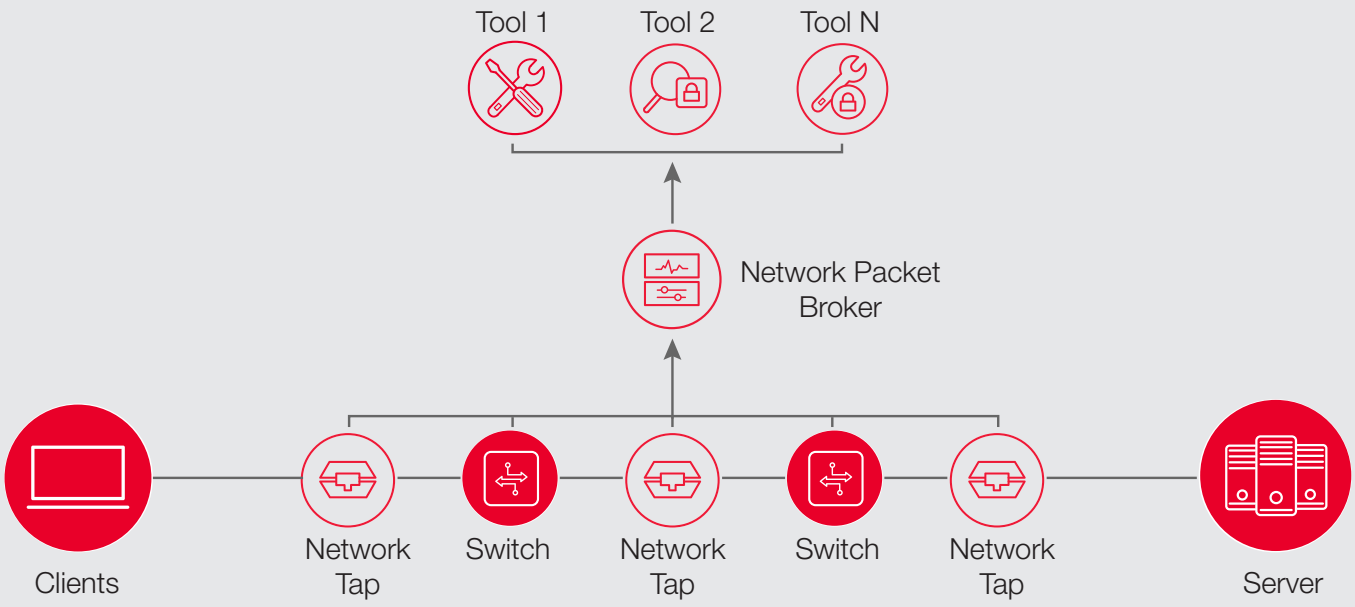

Figure 1: A network packet broker aggregates data from tap and SPAN ports.

By filtering data within the NPB, the monitoring tool is free to perform the work that it was purchased to do, resulting in more useful work being done by the monitoring tool. Since the tools are now as efficient as possible, less tools may be required to accomplish the same goals. In addition, the right choice of an NPB optimizes filter programming costs by removing the manual command line interface (CLI) process used in switched port analyzer (SPAN) ports and some NPB models.

Once a visibility architecture is put in place, you are about two-thirds of the way to your goal. You know have the visibility you need to go with the defensive security tools that you already had.

Security Resilience

Once a visibility architecture is put in place, you are about two-thirds of the way to your goal. You now have the visibility you need to go with the defensive security tools that you already had.

A different, but complimentary, concept is called network security resilience. Resilience is the ability of a system to return to original form, position, etc., after being bent, compressed, or stretched. It is also referred to as the capacity to recover from difficulties. Security resilience is the ability of your security architecture to recover and return to a normal state after an attack and/or breach has occurred. During attacks and breaches, minutes matter. All government agencies need this interval to be as short as possible.

A security resilience approach can provide you the following benefits:

- Strengthen your capabilities to defend against attacks
- Maximize your ability to rebound from an attack
- Minimize the severity and cost of security breaches

Proper data acquisition can bring huge value to your network security systems. When you actually integrate your security architecture with your visibility architecture, you equip yourself with the necessary tools to properly visualize and diagnose the problems on your network.

Once you have the data you need, it can be fed to the security architecture. For instance, in the case of inline deployed security tools, Secure Sockets Layer (SSL) decrypted data can be fed to an intrusion prevention system (IPS), web application firewall (WAF), or data loss prevention (DLP) tool for inspection. Suspicious data can go through further analysis.

Here are some common elements in a security resilience approach:

- Threat intelligence gateways
- Inline tools – FW, IPS, next generation firewall (NGFW), SSL decryption, web application firewall (WAF), etc.
- Network packet brokers – inline and out-of-band
- External inline bypass switches
- Out-of-band tools – intrusion detection system (IDS), DLP, security information and event management (SIEM), TLS/SSL decryption, flow data analysis
- Application filtering, geolocation information, NetFlow data
- Security device testing
- Network penetration testing
- Cyber Range training

If you cannot protect everything 100 percent, then security resilience is the next best approach. It allows you to secure as much of the network as you can while building in network visibility and recovery systems to mitigate the effects of a breach as fast as you can.

If you cannot protect everything 100 percent, then security resilience is the next best approach. It allows you to secure as much of the network as you can while building in network visibility and recovery systems to mitigate the effects of a breach as fast you can.

Mapping Security Resilience to the NIST Cybersecurity Framework

The NIST *Framework for Improving Critical Infrastructure Cybersecurity* provides a common language for understanding, managing, and expressing cybersecurity risk. This framework is built upon concepts to organize information, enable risk management decisions, address threats, and improve through lessons learned. The foundation to these concepts are aligned within five core functions:

- Identify
- Protect
- Detect
- Respond
- Recover

As part of the Presidential Executive Order 13800, all government agency heads will be held accountable for implementing solutions and managing the risks associated with threats to our nation's cybersecurity and thus must take immediate action to review cybersecurity protocols in order to upgrade each department's IT infrastructure. This includes protecting IT systems and data from unauthorized access, detecting and preventing anomalies and incidents, and mitigating impact to affected systems. It is up to the CIO, CISO, and security architects for each government agency to map their processes and security defenses to the NIST framework. Exact actions are obviously agency specific. However, what is not commonly understood is how a visibility architecture can be combined with the security architecture to bolster network security in those five categories.

The following sections provide an overview of how you can combine the two, often distinct, architectures into one coherent architecture.

Identify

This security function focuses on who and what could be affected by various scenarios, governance mandates, and agency risk. The purpose is to gain a better understanding of agency cybersecurity risks associated with specific systems, assets, data, and capabilities. A typical security architecture would focus just on security processes and continuous analysis of relevant security patches along with their priority. However, there are two additional fundamental activities that needs to be addressed:

- Continuous monitoring of critical applications, operating systems, and the network
- Comprehensive monitoring of public and private cloud resources

As the NIST cyber security framework demonstrates, continuous monitoring is important to network security. This is due to the both the visibility (i.e. insight) that is provided into the network and the increase in productivity that is provided. This includes planning for where and how to collect critical sets of data. While the where (edge, core, cloud, etc.) is network dependent, the how is often universal and consists of using taps, NPBs, and various security and monitoring tools.

Virtual taps are needed for public and private cloud instances. These specialized taps allow you access to critical packet data that is still necessary for troubleshooting, compliance, and forensic analysis, even though you may have moved to a hosted cloud environment.

Taps are passive devices that supply a complete copy of the data running on that network segment. Packet brokers are used to filter and refine the data in multiple subsets that one or more tools will need to analyze. Virtual taps are needed for public and private cloud instances. These specialized taps allow you access to critical packet data that is still necessary for troubleshooting, compliance, and forensic analysis, even though you may have moved to a hosted cloud environment. Without this packet data, you will create blind spots within your cloud solution. The integration of these components gives you the network visibility that you will need to properly baseline your network and application performance and then perform ongoing monitoring activities.

Protect

The Protect function outlines how safeguards for the critical infrastructure need to be developed and implemented, supporting the ability to not only limit cybersecurity events, but also activities to contain them. For instance, a typical security architecture would focus on access controls, protective technology (firewalls), safeguarding data, processes, and procedures. Besides the common wisdom just mentioned, there are additional activities that government agencies should consider to increase their defensive capabilities.

These activities include:

- Deploy threat intelligence gateways to eliminate up to 30 percent of threats immediately

- Incorporate the use of security device testing appliances to validate the effectiveness and performance of your security devices before you deploy them in the live network
- Implement cyber range training activities to strengthen the skills of your threat response teams

What is not often considered in typical architectures are threat intelligence gateways that drop traffic to and from known bad IP addresses. These devices reduce analysis load on security tools by eliminating a substantial portion of the threat upfront. One of the drawbacks of a firewall-only approach is that updates to the Internet Protocol (IP) address access lists are typically manual processes. By adding a threat prevention gateway, which uses automated updates to known bad IP addresses, you can significantly minimize both incoming and outgoing traffic to bad actors. A threat intelligence gateway has been shown to reduce the amount of data that needs to be screened by an IPS by up to 30 percent.

Another overlooked security defense is the use of security device testers. These devices provide proactive testing of security tools (firewalls, IPS, web application firewalls, etc.) in a lab environment to ensure that the tools protect and defend the network as advertised.

Another overlooked security defense is the use of security device testers. These devices provide proactive testing of security tools (firewalls, IPS, web application firewalls, etc.) in a lab environment to ensure that the tools protect and defend the network as advertised. It is very common for the actual performance of these devices to be less than manufacturer specifications, because the tools were tested in an ideal lab environment. This is important because If the tools underperform, additional tools will need to be purchased at additional costs. Lack of device testing will lead to holes in security defenses. Processes may need to be adjusted to account for this.

Cyber Range training for IT security personnel is another proactive activity in this area. Cyber threat training gives security personnel the ability to experience real attacks (in a controlled environment) so that they can better recognize the symptoms and precursors on a live network. With this knowledge, security personnel not only respond better to attacks, but they can help the systems recover faster to reduce liabilities.

Detect

The Detection function focuses on continuous end-to-end monitoring to detect anomalies. This component is concerned with identification of actual issues, intrusions, and breaches. This is where a typical security architecture would focus on deploying some set of inline tools (IPS, WAF, etc.) or out-of-band tools (SIEM, DLP, IDS, etc.). In this context, inline refers to being directly in the path of live traffic, i.e. traffic must pass through the device before it continues on. Out-of-band is where data has been copied and sent to devices that are not in the path of the live network traffic and thus do not delay or impede the flow of data across the network.

There are several activities that can be implemented to strengthen the Detect functional area:

- Deploy inline security tools with fail-over technology (bypass & heartbeats)
- Deploy inline packet brokers for improved fail-over scenarios (hot standby, load sharing)
- Deploy SSL decryption for improved threat visibility
- Focus on data loss prevention by filtering monitoring data to out-of-band security tools (SIEM, DLP, IDS, forensic analysis) for faster analysis
- Proactively look for indicators of compromise using application intelligence
- Capture suspicious data with packet captures for either immediate or delayed analysis

When an external bypass switch is combined with an NPB, a very strong and resilient defense can be created. Load balancing, high-availability options, and heartbeat signaling within the NPB can be used to create a self-healing architecture.

Once inside the network, inline security tools are often deployed near the perimeter. This can create a single point of failure, even if the tools have a built-in bypass capability. For instance, what happens when you want to remove the device altogether? An external bypass with heartbeat capability allows you to increase network availability and reliability with microsecond automatic fail-over and fail-back technology.

When an external bypass switch is combined with an NPB, a very strong and resilient defense can be created. Load balancing, high-availability options, and heartbeat signaling within the NPB can be used to create a self-healing architecture. A new threat is the inclusion of malware within encrypted data payloads. This data can be unencrypted by the NPB, or a purpose-built device, and then passed on to security tools (like an IPS or WAF) for further examination. Data that is safe is re-encrypted and sent back to the bypass switch to traverse downstream.

At this point, TLS decryption can also be deployed. Various security and monitoring tools do not support decryption capabilities. If they do, they often consume considerable central processing unit (CPU) resources to perform this function, which ultimately makes the tools less efficient and slower to perform their core analysis functions. The NPB can perform the data decryption before the data is passed on to the various tools. This allows you to decrypt the data only one time and then share across multiple security devices.

Once data moves into the core, it can still be inspected. Instead of a real-time analysis, data can be captured and moved to out-of-band security and monitoring tools for deep packet inspection, forensic analysis, log file analysis, NetFlow analysis, regular expression (Regex) searches, etc. This solution uses NPBs and application intelligence to filter out either Layer 2 through Layer 4 data and send the appropriate data to the corresponding tools, like DLPs, SIEMs, log file collection systems, IDSs, etc.

Once data moves into the core, it can still be inspected. Instead of a real-time analysis, data can be captured and moved to out-of-band security and monitoring tools for deep packet inspection, forensic analysis, log file analysis, NetFlow analysis, regular expression (Regex) searches, etc.

This monitoring data can also be filtered by Layer 7 application type. For instance, maybe you only want to see Facebook data or do not want to analyze Netflix data. That data is emphasized in the filtering process. In addition, the application intelligence gateway can generate NetFlow data and additional data (like geolocation, browser type, device type, etc.) that can be used in further analysis of the data. Data masking, Regex searching, and packet capture (PCAP) capabilities are also provided to help analyze specific data natively within that solution or by third-party tools.

Data captures of application information can also be replayed and analyzed in the lab to make sure the network is working correctly. Common questions include: What exactly triggered the fault? What pre-event traffic was occurring? Is this fix the correct solution? The answer is often packet-based data capture files (which can be created by the NPB) and/or NetFlow-based security data captures which can also be created and viewed in a security lab for tactical analysis of how an intrusion took place.

Respond

The Respond function is about taking action to mitigate the threat and keep it from spreading. Typical activities include taking equipment and/or the network offline, turning off features temporarily, active debugging of the problem/attack, and the coordination and implementation of responses and next actions (informing senior management within the agency, law enforcement, etc.). Some IT departments also focus on using specific tools like a SIEM or DLP to create a faster threat analysis.

There are additional activities that can be conducted such as:

- Deploy automation with representational state transfer (REST) interface to packet brokers
- Continually update threat intelligence gateways to reduce/prevent the exfiltration of data

One of the most powerful, but often overlooked, features for data center automation is automating the out-of-band network monitoring switch. In this case, automation means packet brokers can initiate functions (e.g., apply filters, add connections to more tools, etc.) in response to external commands. This data center automation is akin to software defined network (SDN) capabilities, which allow a switch/controller to make real-time adjustments in response to events or problems within the data network. However, the source of the command doesn't have to be an SDN controller. It could be a network management system (NMS), provisioning system, SIEM, or some other management tool in your network.

Let's look at one quick example. It's 3:00 am and a hacker has just attacked your IT network. How does your network behave? Does it understand who's attacking the network? Once the attack has begun, what exactly happens? Maybe you've purchased a SIEM and just like you hoped, it spots the problem. What will the SIEM do next? Is there an intrusion detection and prevention system that just happens to be connected to the right SPAN port that the SIEM can spin up? What about starting a packet capture? How about starting the forensic recorder? Do those tools just happen to be on the same SPAN port that this threat vector is coming from?

A threat intelligence gateway can be used to prevent data leakage. Since the access list in the threat intelligence gateway is constantly updated, this can be used to augment the IP address filtering process within the firewall, as that is often a manual process. Even if the data initially got through the threat intelligence gateway, the regular update process for the access lists can catch the traffic to/from the now newly identified bad IP address and stop data at the perimeter of the network before it leaves.

Automation means packet brokers can initiate functions (e.g., apply filters, add connections to more tools, etc.) in response to external commands.
This data center automation is akin to software defined network (SDN) capabilities, which allow a switch/controller to make real-time adjustments in response to events or problems within the data network.

Recover

The final function of Recovery is centered on the remediation and repair of the data network and its components. This phase is the main area for repair and restoration of any network capabilities that are damaged. There is also emphasis on prevention of the issue in the future. Typical architecture activities include reprogramming component software, changing passwords, applying patches, and turning off features permanently.

The following additional activities should be considered as well.

- Use security device testers to run "what if" scenarios to validate the efficacy of new security device settings
- Use proactive monitoring to test performance of new features

Once an intrusion or breach has occurred and a solution defined, that solution can then be tested with a security device testing solution. The test solution can stress the equipment and network to its breaking point to see if the fix is a long-term solution or not. This data is critical for network and device dimensioning. As part of the test effort, you can see the real performance impact of decisions with various "what if" simulations (like SSL key and cipher impacts, latency due to SSL, how application intelligence would look with different geographies and traffic mixes, how distributed denial of service (DDOS) mitigation would affect your network performance, etc.). Running these types of simulations is important, because you cannot just cut services to your customers to correct a problem in the long term, you need to restore services but ensure the integrity of your new network configuration.

Once an intrusion or breach has occurred and a solution defined, that solution can then be tested with a security device testing solution. The test solution can stress the equipment and network to its breaking point to see if the fix is a long-term solution or not.

Proactive monitoring uses visibility technology to actively test your network and has several fundamental benefits including the ability to:

- Know immediately the performance level of your network
- Test upgrades in a maintenance window to understand how well your applications are running

Network performance and application performance may sound simple, but these can be difficult to ascertain, especially if certain applications are off-line. To get a true indication of network performance, the network needs to have a large amount of traffic on it, which makes you dependent upon peak busy hours. This solution allows you to place probes anywhere in your network and test whenever you want to. It also allows you to accurately simulate the right traffic with synthetic transaction so that application performance management (APM) tools can observe how well applications are truly performing. For instance, this allows you to simulate small packets or Skype-like data if you want to test your instant message (IM)/voice/video solution. Because of the synthetic traffic generator, this can be performed within the maintenance window (or anytime) to lessen the impact to network users.

Conclusion

New initiatives for government networks, like the NIST *Framework for Improving Critical Infrastructure Cybersecurity*, require an adjustment to common network management processes. As part of the new standard, government networks need to operate within a context that supports an Identify, Protect, Detect, Respond, and Recover approach. This includes all of the network, not just security-focused appliances and solution. As part of this initiative, a visibility architecture should be created that not only focuses on capturing the requisite data, but supports the security architecture in the goal of preventing security breaches.

Here are some specific things that you can implement to strengthen network security and monitoring practices:

- Use continuous monitoring of critical applications, operating systems, and the network using taps, NPBs, and purpose-built tools
- Deploy threat intelligence gateways to eliminate up to 30 percent of threats immediately
- Incorporate the use of security device testing appliances to validate the effectiveness and performance of your security devices before you deploy them in the live network
- Implement cyber range training activities to strengthen the skills of your threat response teams
- Deploy inline security tools with fail-over technology (bypass & heartbeats)
- Deploy inline packet brokers for improved fail-over scenarios (hot standby, load sharing)
- Deploy TLS decryption for improved threat visibility

- Focus on data loss prevention by filtering monitoring data to out-of-band security tools (SIEM, DLP, IDS, forensic analysis) for faster analysis
- Proactively look for indicators of compromise using application intelligence
- Capture suspicious data with packet captures for either immediate or delayed analysis
- Deploy automation with REST interface to packet brokers
- Continually update threat intelligence gateways to reduce/prevent the exfiltration of data
- Use security device testers to run “what if” scenarios to validate the efficacy of new security device settings
- Use proactive monitoring to test the performance of new features

Keysight network visibility solutions are a powerful way to optimize your network monitoring architecture and strengthen your network security. For more information on network monitoring solutions, visit www.keysight.com/solutions/network-visibility

Learn more at: www.keysight.com

For more information on Keysight Technologies’ products, applications or services, please contact your local Keysight office. The complete list is available at: www.keysight.com/find/contactus

KEYSIGHT TECHNOLOGIES