


How to Secure Your Network...

Through Its Lifecycle



Keith Bromley and Kathy O'Neil. First Edition.

ixia

Contents

Executive Summary	4
Chapter 1: A Shift in Network Security Protection Thinking is Required	7
It's Time for a Fundamental Mind Shift	11
A Systems Approach that Follows Network Lifecycle	13
Chapter 2: The Lifecycle Approach to Network Security	15
Component 1: Controlling Access	19
Component 2: Creating and Enforcing Security Policies	20
Component 3: Perform Network Security Architecture Assessments	21
Component 4: Implement A Well-Planned Network Monitoring Solution	23
Chapter 3: Pre-deployment Security Analysis Validates Your Design	29
Start with the Right Security and Visibility Technology	32
Optimize Performance of Components as They Relate to the Whole System	33
Validate Network and Policies with Real Malware	35
Validate Against Known Threats	36
Chapter 4: Installation, Testing, and Commissioning Ensures Design Compliance	39
One More Check that You Have the Right Systems in Place	40
Baselines Are Critical for Fallback, Alerts, Change Management, and Ongoing Assessment	42

Chapter 5: Production Network Visibility Enables Real-Time Vigilance	45
Finding The Perfect Solution: Combining Visibility and Security	46
Implementing In-Line Visibility Solutions	48
Extend Physical Network Monitoring Tools into the Virtual Realm with Virtual Taps	52
Implementing Out-of-Band Visibility Solutions	55
Visibility Architecture Automation to Improve Security	58
NPB Application Intelligence Improves Security	63
Chapter 6: Software and Hardware Upgrade Testing Limits Security Flaw Introduction	67
Upgrades – Where Dangers Lurk	69
3 Steps to Reduce Change Management Risks	71
Best Practices for Secure Change Management	75
Chapter 7: Conclusion	77

Executive Summary


As daily news stories continue to document, enterprises are struggling to prevent breaches to their networks. Not only do breaches affect the company brand, but economic losses continue to mount as well. There needs to be a fundamental mind shift away from the current thought that network security is a one-time thing or a one-size-fits-all. To create a resilient system, network security needs to be an ongoing process, not just occasional technology implementations. If not, enterprises risk exorbitant remediation costs that may threaten company viability.

Security professionals should look at a lifecycle approach to securing their networks. Just as the seasons change during the course of a year, your network has different needs and attributes during the four main lifecycle stages of its usable life. Understanding these four lifecycle stages and their different security needs will minimize the threat to your network.

The benefits of this lifecycle approach can be summarized as follows:

1. **Pre-deployment** security testing validates your security design
2. Threat assessments during the **installation and commissioning stage** ensures design compliance
3. Inserting a visibility architecture into the **production network** enables real-time visibility and security vigilance
4. **Software and hardware upgrade** testing minimizes and/or eliminates security flaw introduction

This book will present the four stages of your network lifecycle and the best practices to securing them. Consistency of monitoring and security policies will be required to make sure that the desired benefits are actually achieved.



Chapter 1:

A Shift in Network Security Protection Thinking is Required

The security threat to today's enterprise networks is imminent. It's not a question of if, but when you will be attacked and how well will you survive that attack. And then, how will you know you are/were attacked? With about 25 percent of companies seeing more than 21 attacks per month ([Arbor Networks' Worldwide Infrastructure Security Report Volume IX](#)), you have most likely already been attacked. If you have a cloud environment, then the [Ponemon 2014 Data Breach: The Cloud Multiplier Effect report](#) predicts that you are 3 times more likely to be breached.


So, are you positioned well enough to counter the onslaught? Many have not been – as you can easily see in the news reports of the recent high-profile security breaches: Facebook, Twitter, Apple iCloud, Microsoft, Target, Neiman Marcus, Michaels, Sprouts, Basha's, Adobe, Home Depot, Jimmy John's Restaurants, Sears, Kmart, Dairy Queen, JPMorgan Chase, Staples, Marquette University, UC Davis Health System, Oregon Employment Department, North Dakota State College, Sally Beauty Supply, SuperValu, etc. These are all well-known entities that have had their brand tarnished and many have paid a heavy price within the last year or so.

As Symantec notes in their [Internet Security Threat Report for 2014](#), 2013 was the “Year of the Mega Breach”. The total number of breaches was 62% higher than the level set in 2011 (which Symantec had previously cited as the “Year of the Data Breach”). In addition, the number of exposed records was much higher in 2013. In 2011, there were five reported security breaches with more than 10 million records exposed for each breach. In 2013, that number

was eight with breaches resulting in more than 552 million identities exposed. The Online Trust Alliance (OTA) saw similar data and noted in their [2014 Data Protection & Breach Readiness Guide](#) that 40% of the largest recorded security breaches occurred in 2013.

But what does this mean? Is this just an aberration? With trend data indicating that the number of attacks and breaches continuing to increase, it's doubtful this is an aberration. Rather, it appears to be a systemic problem. This is especially true when you consider the root cause of security breaches. According to the [Ponemon Institute's 2014 Cost of Data Breach Study: Global Analysis](#), only 42% of security breaches are caused by malicious and criminal attacks. The other approximate 58% of breaches were due to human error (30%) and system glitches in IT and business processes (approx. 28% depending upon rounding).

Data Breach Root Cause Analysis (for 2013)




Source: Ponemon Institute 2014 Cost of Data Breach Study: Global Analysis Report, May 2014

A [Forrester research survey for 2013](#) also found similar results. The Forrester study indicated that 36% of breaches were caused by inadvertent insider misuse while another 25% of breaches were

caused by malicious insider actions. One note, the survey allowed for multiple answers since multiple factors can enter into the root cause of a breach.

“What were the most common ways in which the breach(es) occurred in the past 12 months?”



Base: 512 North American and European enterprise and SMB IT security decision-makers whose organizations had a data breach in the past 12 months

Source: Understand The State Of Data Security And Privacy: 2013 To 2014, Forrester Research, Inc., October 1, 2013.

Separately, the Open Security Foundation also found in their report, [An Executive’s Guide to 2013 Data Breach Trends](#), that approximately 31% of breaches were due to lack of internal controls or employee fault in 2013. The root cause of these problems is usually a lack of defined and/or proper policies for users and lack of defined or proper procedures for information systems. Fixing these two areas, and enforcing them, addresses the systemic portion of the 58% of breaches just mentioned.

According to the [2014 Trustwave Global Security Report](#), breaches were focused in three main areas: e-commerce (54%), point of sale (33%), and data centers (10%). However, if you think it's all about payment card data theft, then you would be wrong. Trustwave also determined that 45% of the data stolen was about other personal identifiable information (PII) like credentials, internal communications, and other personal information. Ponemon also found in a report they wrote on healthcare privacy in late 2012 that healthcare fraud was on the rise (a 52% increase over the previous year) whereby criminals were using the PII of legitimate people to get illegal medical care or to defraud the healthcare industry. In fact, stolen healthcare identities can sell for \$18 a piece versus about \$1 for stolen credit card accounts.

One more data point reported for last year is that intruder attack methods are somewhat diverse. [Verizon's 2014 Data Breach Investigations Report](#) (DBIR) found 9 fundamental threat vectors. This means that you need to look at the whole network to protect it: architectures, networks, endpoints, applications, databases. Point solutions aren't going to be able to address everything – it's going to take a coordinated effort to create a truly secure and resilient architecture.

It's Time for a Fundamental Mind Shift

So, again, what does this mean? The logical conclusion is that what's being commonly implemented for a security architecture doesn't really work well. There isn't a product (or two) that can be bought and placed into the network that will magically fix it. So, there needs to be a fundamental mind shift that network security isn't just something you go out and purchase, but something that

you work on every day. At the same time, you can optimize network security by using data-driven methodologies to objectively analyze network components and the network architecture. The secret is in the mixture of policies and products to create the correct proactive and reactive security architecture for your business. In addition, the security architecture MUST include a full visibility plane that allows IT clear visibility into the network to “see” hidden threats and to collect proof and depth information about breaches. The visibility plane is also a must for ALL forms of lawful intercept.


The second conclusion is that you need to invest in creating and optimizing your security architecture now, while the cost is essentially “pennies on the dollar” when compared to the remediation costs for a breach. Those remediation costs can become exorbitant, as we have seen play out for most of the companies we just cited that were breached in the last year. Depending upon company size, remediation costs can even threaten the viability of the business. In addition to protecting your intellectual property (IP), which is typically the most valuable piece of the company, government fines and civil lawsuits can cost millions of dollars.

For instance, YourTel America and TerraCom were fined \$10M on Oct. 27, 2014 by the FCC for intentionally not protecting their user data. On May 7, 2014, the U.S. Department of Health and Human Services Office for Civil Rights completed a HIPAA investigation of New York Presbyterian Hospital and Columbia University. Both were found guilty of releasing patient information inappropriately and were fined \$4.8M for not protecting their patient data.

A Systems Approach that Follows Network Lifecycle

This is where taking more of a systems approach that follows the lifecycle of your network can improve its resiliency in the face of all the different security threat vectors that exist. The meaning of lifecycle here (not to be confused with other IT processes like the Deming cycle where you have the Plan, Do, Check, and Act phases) is that any business network goes through different phases during its usable life. These lifecycle phases are:

- ▶ Pre-deployment (before it's introduced)
- ▶ Installation and commissioning (initial deployment)
- ▶ Production (steady state)
- ▶ Upgrades (change management for hardware/software upgrades)



Each of the lifecycle stages has distinct attributes and capabilities that need to be assessed for vulnerability. As you build your security architecture, breaking down your policies and activities to match the lifecycle of the network will make your processes simpler and more repeatable.


Chapter 2: The Lifecycle Approach to Network Security

A key point to understanding a lifecycle approach is to understand that it augments your existing security and visibility architectures. The change is that you are simply analyzing your network at different points in time – from its inception to deployment, and then various instances later on as it is modified and matures. Everything, including a network, has a lifespan that is full of changes throughout the course of its life. It's fundamentally necessary to observe this and implement the requisite changes at the right time.

As the root cause data from Chapter 1 shows, protecting network security is not just about external attacks. This is further compounded by data from the 2014 Trustwave Global Security Report that shows that 71% of compromised victims did not detect the breach themselves. In addition, the median number of days from initial intrusion to detection was 87. When you combine this with the data above that 58% of breaches are due to “inside” threats, there has to be another (systemic) element involved in the network security picture.

This is where a “systems view” to network security might be a better way. This thought process allows you to step back and take a comprehensive view to network security. Instead of focusing on just the microscopic view of the physical, people, and network security aspects, the systems view allows you to take a macroscopic viewpoint and consider the following items:

- ▶ Alignment of company (physical and data) security with business objectives
- ▶ Risk-based analysis of security threats

- 
- ▶ Balance among people, processes, and technology in your security processes
 - ▶ Convergence of multiple security strategies

Before we dive further into the lifecycle of network security, let's take a quick look at the fundamental security architecture. You'll want to understand your network before attacks happen – what areas are strong and what areas are weak? In short, a good network security architecture will consist of multiple layers of security. You need to control:

- ▶ Access
- ▶ Policies and procedures
- ▶ Architecture performance
- ▶ Continual auditing/monitoring

Better security architectures will also have both proactive and reactive components. Proactive capabilities focus on vulnerability identification and then remediation of those threat vectors while reactive capabilities address anomaly detection and resolution. One approach or the other probably won't protect you from advanced hackers. To really understand your network, you need to characterize it first. While no one can understand everything, you can definitely create and record a baseline of how your network behaves so that you can look for patterns in normal and abnormal behavior. This will allow you to see the attack and quickly scope the impact to your network.

Let's look at the four fundamental components in a traditional network security architecture.



Component 1: Controlling Access

The first fundamental area to consider is access, which includes the basic building blocks you've already been using to control access: firewalls, intrusion detection system (IDS) and intrusion prevention system (IPS), and specialized security monitoring tools. These are all standard tools. The firewall is used to block initial entry. Once a potential threat gets past the firewall, it needs to pass an inspection by an IDS. The IDS monitors networks and systems for abnormalities or policy violations, but it doesn't block traffic. This is the job for the IPS, which can then drop packets, block traffic from specified IP addresses, reset connections, etc.

There are several sites like [InfraGard](#) that provide a list of bad sites that should be always added to the firewall. Additional sites like [CVE](#) and the [National Vulnerability Database](#) can help as well. Professional hackers will attack the routers, switches, IPSs, and firewalls as primary targets, so you need to protect these devices.

After the anomaly has run the gauntlet, you can still capture data about it through specialized security monitoring tools. For instance, data recorders can record all traffic, or certain portions of it, depending upon your threat level. Other tools can investigate patterns, trends, and policy breaches for further data analysis. You can also use this forensic data to analyze threat vectors (multi-vector attacks, advanced persistent threats, etc.) to prevent future attacks. This isn't the best-case scenario, as you'd rather prevent all attacks. But the reality is that you probably can't, so you need the ability to analyze the successful attack for root cause information.

After these initial layers of security, you can still add more security like: next-generation firewalls (NGFWs), distributed denial of service (DDoS) mitigation, a sandbox for advanced threat detection, anti-malware, honeypots, etc. Unfortunately, it may not mean better security. While they will provide incremental improvements in security, you need to integrate these tools and capabilities in a planned way with your basic security architecture to make sure they will actually work as desired and deliver maximum value for the additional costs.

Component 2: Creating and Enforcing Security Policies

The second fundamental area is to create the appropriate policies. This includes:

- ▶ Usage for what is allowed and forbidden (e.g., network, company assets, code of conduct)
- ▶ Escalation
- ▶ Reporting



Policies are key to dictating how the network can and cannot be used by personnel, along with the process for countering security attacks and escalating information about them. Reporting is a key aspect. Many companies have been burned in this area, so having a clear, written policy of how and when security attacks and breaches should be documented is extremely important. Another aspect of reporting should involve the periodic evaluation and documenting of security tool performance. For instance, do any of the security tools have log files showing attacks, loss of data, etc. and how often is this data reviewed?

While having policies in place are a good practice for network security, they will be critical to executive management if a breach does occur because the executive team (and company) will want to be able to show due diligence in securing company assets and intellectual property.

Component 3: Perform Network Security Architecture Assessments

The third area is to assess the network so that you understand the following about it:

- ▶ Security effectiveness
- ▶ Performance under load
- ▶ Product vulnerabilities
- ▶ Visualization of network limits

There are several common security assessment options for network assessment including the following activities:

- ▶ Penetration testing
- ▶ Vulnerability assessment
- ▶ Security audits
- ▶ Code testing

However, these methods have drawbacks like:

- ▶ Assessments use static glimpses of the network
- ▶ Tests are often limited in scope
- ▶ They fail to stress test components and the system
- ▶ They don't mimic the real world

To reduce security uncertainty, you need more comprehensive security assurance and validation that involves the following aspects:

- ▶ Test with real-world, recent attacks
- ▶ Use line-rate application loads
- ▶ Assess devices and systems
- ▶ Validate data sheet performance claims of your security components
- ▶ Scale testing to cover extensive test cases

Component 4: Implement A Well-Planned Network Monitoring Solution

After you've finished your initial audits, you'll want to include methodologies to perform continued monitoring of the network. When in the monitoring phase of your security architecture, you'll want to create an architecture that focuses on three core aspects: prevention, detection, and response. Jim MacLeod talks further about these in his blog ([Role of Packet Capture in Network Security](#)) posted March 7, 2013 on www.lovermytool.com.

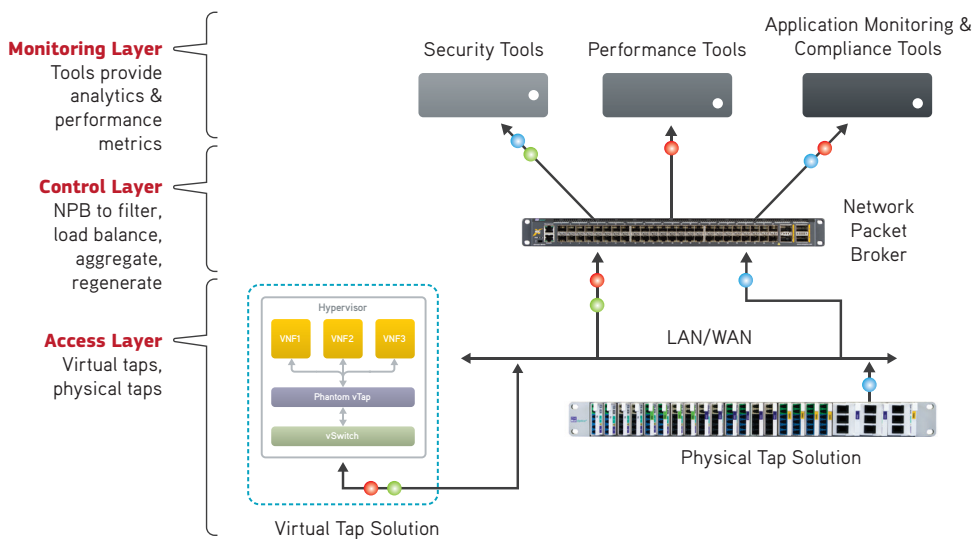
However, just adding monitoring tools to a security architecture isn't enough. There are still many blind spots and inefficiencies that will appear in your network. This is where you also need to have a visibility architecture that focuses on allowing you access to the right data at the right time that can be forwarded to the specialized monitoring tools for proper analysis.

It's very important to set up a visibility architecture correctly. More information on this subject is available in the Ixia whitepaper "[Creating a Visibility Architecture](#)", but the gist is that experience shows that the best way to increase network visibility, optimize network performance, and improve network security is to step back and take a holistic view of where you are and where you need to be, and then create an achievable path to get there by accounting for all necessary monitoring and visibility components.

Visibility access is a very important piece of a visibility and security architecture. This is the portion of the visibility architecture that handles how you tap into the network. For instance, some people

are still using SPANs or VACLs, while others are using taps. In either situation, you need the best solution possible to enable your security architecture.

After the access layer, you'll need data control capabilities to properly segment and de-duplicate your monitoring traffic. Finally, there are the monitoring tools that will analyze and distill the data into coherent information. A well-designed visibility architecture will then enable the security architecture by supplying the proper data it needs at the right time. An example is shown in the following diagram.



In regards to the access layer, SPANs, VACLs and taps can all be used. However, taps have become the preferred method. In fact, most international networking standards are now specifying the use of taps. The reason is simple. By their nature, SPANs and VACLs are generating summarized data that may have delays associated with the data packets. In addition, SPAN ports tend to create a lot of duplicate data on the network. This duplicate data

has to be analyzed and removed which wastes monitoring tool and packet broker CPU cycles and decreases the monitoring tool efficiency.

It should be noted that SPAN ports and VACLs are very vulnerable to attacks and hacks due to command line interface capabilities that can be exploited. In contrast, most taps don't have this issue. This means that the SPAN ports and VACLs can actually compromise any security visualization process. They also cause issues like dropped packets (both good and bad), create duplicate packets and create timing issues. Also, with distributed attacks, these devices can become overloaded where important attack information is lost, especially in multi-vector attacks where some of the attacks will be seen while others will be missed. This compromises the ability to see the full threat profile.

Taps come in different forms so you'll need to figure out the right mix for your network. Typical tap categories include in-line tap, traditional hardware-based tap, and virtual tap. Most networks will use several kinds of taps as well as several taps of each kind.

Network packet brokers (NPBs) form the basis of the control layer. These are the middlemen that remove unwanted/unnecessary data (either through filtering or deduplication). Other functions include load balancing data across tools, data aggregation, packet slicing, time stamping, and port tagging. A few of these devices can go a long way to improving network visibility, especially if extended functionality like application intelligence and filtering are included.

Monitoring tools are the third layer. These types of tools vary depending upon the visibility needs. Functionality can include simple monitoring, deep packet inspection (DPI), troubleshooting, regulatory compliance, lawful intercept, and security defense and protection. Typical vendors include: FireEye, HP, CA, Riverbed, and Fortinet.

These three layers form the basis of an architecture that can give you two-way visibility into what's happening on the network and then feed the requisite data to the security architecture for resolution.

Now that we've looked at a traditional design approach for a network security architecture, there is a better way. Consider adding a network lifecycle approach to securing the network.

Just as the seasons change during the course of a year, your network has different needs and attributes during the four main lifecycle stages of its usable life. Understanding these four lifecycle stages and their different security needs will minimize the threat to your network.

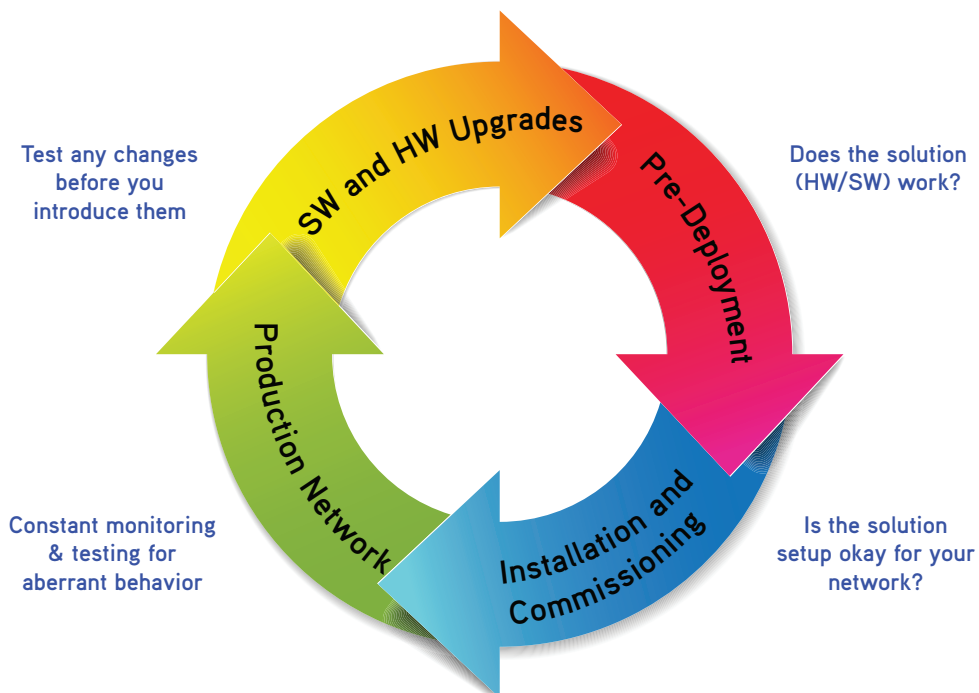
The benefits of this lifecycle approach can be summarized as follows:

- ▶ Pre-deployment security analysis validates your security design
- ▶ Threat assessments during the installation & commissioning stage ensures design compliance
- ▶ Production network visibility enables real-time security vigilance

- Software and hardware upgrade testing prevents security flaw introduction

Here is the basic network lifecycle. These are the four stages that we are concerned with. There is a fifth stage, End of Life, that is not discussed here for obvious reasons. Each of these stages has unique characteristics that require different security analysis tools and methodologies to ensure optimal performance.

As shown below, each stage also has a fundamental goal that is associated with it.

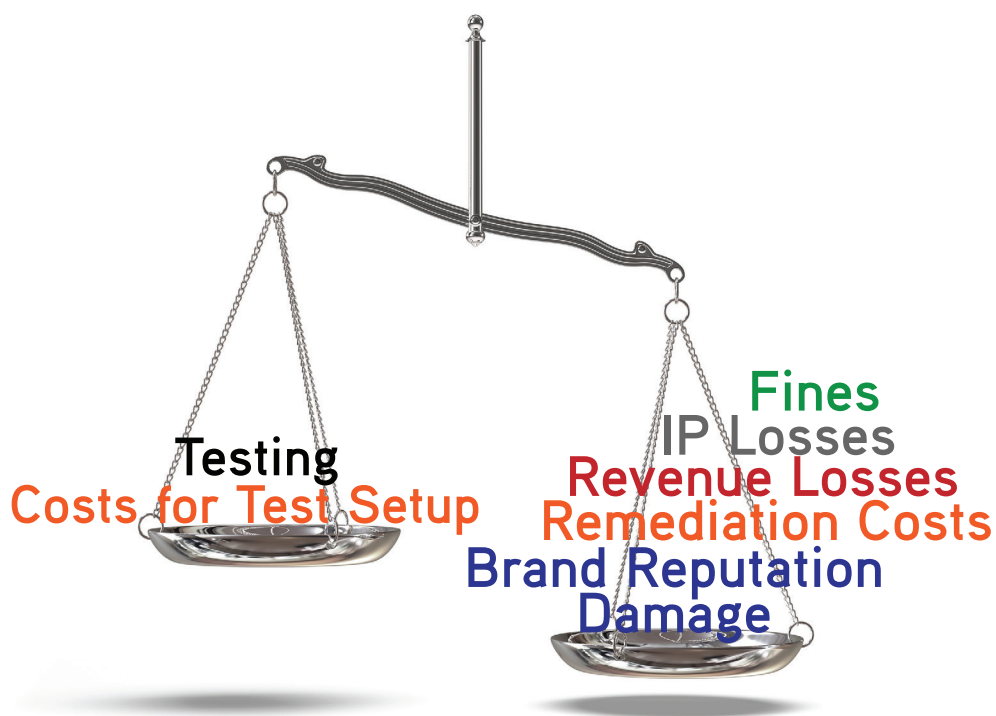


The next several chapters will perform a deep dive into each of the lifecycle stages and show how to achieve the fundamental goal for each stage.

Chapter 3:

Pre-deployment Security Analysis Validates Your Design

One of the most important things you can do is test your network, applications, and access methods. Some in IT hate the concept of testing their networks because of the time and cost involved. What they don't understand is that their network IS going to be tested. Either they test it, or someone else (of a less trustworthy nature) will. The former method may cost some time and effort, but the latter method will cost at least an order of magnitude more of both and will probably cost them their job and a negative employment reference as well.



Unfortunately, pre-production validation can't be left to others. As an example, Apple Inc. recently released their new iOS 8, which had passed its product release testing and processes. However, once the operating system hit the market, it had several zero-day

problems that affected users (see the Computer Weekly article [Apple rushes out second iOS 8 update](#) for more information). Some of these problems were related to IPv6 support (as any Google search for “problems with iOS 8 IPv6” will show you), while others were related to financial environments that resulted in lots of problems with banking applications not working with Apple products. See a good summary on the Middlesex Bank website titled [Online Banking and eDeposit Functionality Problems for iOS 8 Operating System Users](#). When network security and your job are on the line, you can’t afford to hope that the product vendor has done all of the relevant and necessary testing to make sure it will work in your environment.

Another example is the introduction of IPv6. Implementations of this definitely need to be validated before you introduce it into your network. This new protocol has many security benefits, but also has security threats associated with it (see the [Dark Reading](#) website and [Study of IPv6 Security Vulnerabilities](#)) including:

- ▶ **Denial of service** – the sender can add extra headers that can be exploited as part of a DoS attack.
- ▶ **Trespassing/Type 0 routing header** – the network discovery capability lets someone change the path for packets to travel, which means that they can go to places they shouldn’t.
- ▶ **Attacks due to auto-configuration** – this can exploit the duplicate address detection (DAD) feature of IPv6 to allow an attacker to reply to the DAD requests to create a DoS attack.
- ▶ **Filtering device bypass** – if an older device (e.g., firewall) doesn’t understand IPv6 then threats could become hidden as it passes the packets along.

- ▶ **Anycast becomes unsafe** – the header 0 function of IPv6 can be used to allow multiple entities to process Anycast capabilities and form man-in-the middle attacks.
- ▶ **IPv4 to IPv6 dual stack security risk** – tunneling and dual stacks can lead to security issues, especially for IPv4 components. Consider using one or the other but not both.

Start with the Right Security and Visibility Technology

Conducting a data-driven proof of concept in the pre-deployment phase is critical to adopting the right technology for your network and its defenses. This includes evaluating all aspects of your network to anticipate as many issues as possible. While not everything will be accounted for, catching 95 percent of the problems up front before the solution goes live is much better than catching 25 percent. A proper pre-deployment testing approach should evaluate the four following aspects of your security and visibility architecture:

- ▶ Evaluate the actual performance for all security products individually and then compare to vendor specification sheets
- ▶ Optimize network configurations based upon the actual product performance
- ▶ Conduct pre-deployment testing to validate your solution against known threats
- ▶ Conduct pre-deployment testing to validate your solution against known attack vectors

It is a well-known fact that security vendor device data sheets can be misleading. Not because of deliberate misstatements, but because “best case” data is presented instead of “real world” data. To get to the real truth, companies need to put security products to the test to see how they actually perform, preferably before the devices are put into the live network.

The device evaluation testing suggested here can yield real-world performance data to help you determine in a short amount of time which device is best suited for your specific needs. All of your security devices should be evaluated to make sure that they function as specified. This includes: firewalls, IPS, DLP, etc. Anything that inserts to the network can become a problem.

Optimize Performance of Components as They Relate to the Whole System

A second pillar of your security architecture should then include optimizing your network based upon component assessment results. If a device doesn't function as advertised, then it needs to be replaced or the security architecture modified to diminish, if not completely eliminate, the potential threat. The key point is to evaluate your solution once you've selected all of your components to determine the most efficient configuration that allows optimum overall performance.

A secondary benefit can be that if the equipment you've evaluated doesn't work as described, negotiate a discounted price based upon the real-world test results. This is especially true if the device is critical to your architecture. Since it doesn't work as specified,

you're probably going to have to rework the security architecture design, which is going to cost you extra money so you might as well see if you can recoup some of that expense.

CASE STUDY #1

Online Entertainment/Media Corporation (Product & Network Evaluation)

BACKGROUND CUSTOMER is an online entertainment company and has very high performance requirements for their data centers due to massive growth of inline applications. They had three main goals:

- ▶ Assess architecture performance
- ▶ Expose limitations
- ▶ Evaluate customer QoE

SETUP An identical copy of CUSTOMER's network was setup in a pre-production staging environment. In this pre-production environment that included a NGFW firewall, load balancer with security module, and an IPS.

RESULTS This was a pre-production validation of latency related to the delivery of their online content. CUSTOMER needed to keep the real-time transport protocol latency to under 5ms. As they started increasing the connections per second, the latency started to drop.

CUSTOMER found major flaws in the design of the network. Particularly, they identified a bottleneck for performance scalability within their current firewalls. After rigorous testing and re-design, the network configuration was updated to maximize the total connection capacity and connection rates.

CUSTOMER observed how the NGFW would react during the actual attack – how it will report the attack in real time and apply mitigation measures, close connections, etc.

The Next Gen firewall needed to maintain a low latency during the failover. The advertised conditions on the data sheet did not correspond to real test results.

BENEFITS

- ▶ The testing revealed a 70% performance impact to the security features
- ▶ The potential fault was isolated and remedied
- ▶ A different vendor was selected for the firewall based upon the objective test results

CASE STUDY #2

Financial Services Company (Product Evaluation & Network Optimization)

BACKGROUND CUSTOMER needed to scale and re-design their network infrastructure to accommodate application needs for mobile payment systems. The first question was to understand how the redesign affected their network security. Then CUSTOMER wanted to understand if they could replace their incumbent IPS devices with next-generation firewalls and not be blind-sided by security threats. Finally, CUSTOMER wanted to collapse their web application firewall into the load balancer solution they already had. Their goals can be summarized as follows:

- ▶ Evaluate all design options
- ▶ Certify devices
- ▶ Validate security architecture

SETUP Their network configuration was reproduced in the pre-production network and evaluated using a security threat assessment tool.

RESULTS Their next-generation firewalls were found to have security issues for the primary configuration. This configuration was changed to eliminate any security architecture holes. An appropriate DDoS detection and prevention tool was found out of four vendor devices tested.

BENEFITS

- ▶ Re-design of the security architecture due to NGFW limitations
- ▶ Restructure of the Application Security architecture
- ▶ Best DDoS appliance selected

Validate Network and Policies with Real Malware

The third key aspect of network security analysis in the pre-deployment phase is to validate your network against malware attacks. This includes viruses, Trojans, worms, hidden applications, rogue users, etc. Banging away at the network like you've got a sledge hammer is the only way to know how it will truly react. Everything else is just a collection of theories.

Besides validating the network against current malware, you'll need to evaluate your security policies that you created to prevent

future malware events. For instance, does your current security policy state a primary method for inspecting email attachments for malware? Does the policy actual work? How could someone get around your policy and how do you intend to validate that policy? All it takes is for one user to get around the policy and they can exfiltrate data or introduce applications that contain some piece of malware that inserts itself into the network.

Validate Against Known Threats

The last fundamental activity in the pre-production stage for your network is to validate it against known threat vectors. You'll need to test against DoS, DDoS, port scans, Trojans, worms, advanced persistent threats and brute force attacks other network attack vectors.

These are the common "front door" attacks. According to the Ponemon Institute's 2014 Cost of Breach Study, approximately 30% of attacks are from malicious and criminal sources and approximately 3% of the total attacks are denial of service attacks.

CASE STUDY #3

Financial Trading Company (Threat Vector Assessment)

BACKGROUND CUSTOMER is a progressive financial trading group and wanted to see how their network would perform in a real-world security attack before it actually happened. The main goal was to stop guessing what would happen and actually measure the results. They were particularly concerned about DDoS attacks and needed to understand the resiliency of their DDoS mitigation service provider. Their three main goals can be summarized as follows:

- ▶ Evaluate DDoS service provider
- ▶ Assess impact to legitimate traffic
- ▶ Evaluate workflow

SETUP A copy of CUSTOMER's network was simulated in a pre-production staging environment. In this pre-production environment, the servers were subjected to volumetric DDoS attacks by an attack generator. As part of the evaluation, legitimate traffic needed to be allowed through while CUSTOMER's network still defended against legitimate security threats, i.e., the security defenses were not allowed to simply "block all traffic" automatically.

RESULTS The setup mimicked service delivery and the impact that this type of threat vector would have on their production network. There was a real flaw exposed in the data center in how the workflow of DDoS attack defenses responded to the waves of attack. The DDoS attack caused an internal network infrastructure failure within the service provider network that caused the US Federal Bureau of Investigation (FBI) to get involved.

BENEFITS

- ▶ CUSTOMER's initial network architecture would have resulted in a breach
- ▶ The potential fault was isolated and remedied
- ▶ Additional systems & processes were further optimized as part of the evaluation

Chapter 4: Installation, Testing, and Commissioning Ensures Design Compliance

Once you've validated your network architecture in the pre-production lab, the next step is to introduce it into your production network and validate it there as well. You'll want a baseline of the network's performance so that you know exactly how it functions "normally" in the live network. This will become a critical component of your future security evaluations as you'll need this baseline to evaluate if any hidden threats have infiltrated the network or if IT personnel have made any unapproved changes to the network. You'll also use the baseline as a rollback position for deployment failures and catastrophic network reconstruction from a security breach.

One More Check that You Have the Right Systems in Place

Part of this process is to validate the security and visibility architecture in a real setting. Now that everything is in place, do you actually have the right solution for the right need? Was anything missed in the initial design? Or did the network change during the project? For instance, if you need to haul heavy loads of construction equipment with a personal vehicle, you're probably going to use a pickup truck, not a Mini Cooper – even though the original budget specified using a Mini. It's the same with your network. Does the design actually function as you need it to, or does it require modification? This isn't a given. According to a [McKinsey & Company study](#), 17% of IT projects fail so bad they threaten the existence of the company and another 56% deliver less value than expected. [According to the Whir](#), 63% of cloud deployments fail the first time they are launched.

Network and data security is not a simple job. Attacks can come from the inside (users), direct outside attacks, or from software, including upgrades. The network manager and engineers must stay aware of all the latest attacks, including component threats and attack methods from social engineering to covert mechanisms. This must all be part of your testing and commissioning procedures. As mentioned before, some common sites for the latest threats and vulnerabilities include: [InfraGard](#), [Common Vulnerabilities and Exposures database](#) and the [National Vulnerability Database](#).

This lifecycle stage includes everything from cabling, installation, and physical access through to the applications and business criteria, and programming – so it should be looked at and implemented with all of this in mind. Wireless access points should be handled with all the diligence of all other testing and commissioning assurance policies. In fact, wireless access security should have its own version of all policies and procedures.

All of the test and commissioning procedures you create will come in handy when your corporation applies for cyber insurance. Your level of documentation for policies and procedures (or lack thereof) can directly affect cyber insurance policy costs.

Commissioning is the final step in deploying the network and giving access to the users, so everything from the physical layer through procedures and policies must be tested for continuity and usability. This includes all emergency procedures for power, mitigation, etc.

Baselines Are Critical for Fallback, Alerts, Change Management, and Ongoing Assessment

Remember, this is the fallback position for all future changes and modifications so it must be as close to bullet proof as possible. All the baseline data will be used to recognize aberrant behavior, attacks, persistent threats, network issues, etc., so it is a very valuable exercise and every effort should be made in recording as many variables as possible for future comparison. These measurements and comparisons will likewise be used to establish your alarm and alert conditions for your network management system (NMS) and/or management monitor and status indication. These records will also allow for stability and efficiency comparison for all future changes or events.

This stage is also where you can set up a practice to continually validate your solution with proactive maintenance window scans and assessments. Since this is now your live network, you won't have an opportunity to evaluate the whole system at one time but, with the right tools, you can evaluate separate portions of your network individually to create a report that can be compared back to the baseline profile.

A separate opportunity exists here to fully understand the capacity and the breaking point of your network. Again, you will only be able to validate certain sections at a time (in the maintenance window) to determine how the network truly performs. These performance assessments are important to ensure resiliency against security attacks.

It's one thing to say you want to protect yourself from denial of service attacks. But which ones have you run assessments against? Is it a traditional buffer overflow attack, distributed DoS, Tear Drop (using weird shaped packet flows), or was it something else? Or, was there no testing at all?

Is a denial of service attack even your real threat? Some of these are shadow attacks, where the DoS attack is a misdirection (sometimes called sliding, vectoring, or cabbaging) and is intended to distract you while malware is inserted or activated on your network to do the real damage.

So, is your network ready for battle? Because it **will** be put to the test.

Chapter 5: Production Network Visibility Enables Real-Time Vigilance

Any production network will require some level of security monitoring. Even if it isn't deemed to be at high risk, the network probably isn't flawless so there are always opportunities for security holes to be introduced or the security on network segments to be marginalized. A simple case in point is the BlackEnergy malware attack. The U.S. Department of Homeland Security announced Oct. 29, 2014 that a variant of a Trojan horse malware program called BlackEnergy had been found in portions of the American critical energy infrastructure.

What's more concerning, the malware didn't appear to be recent. Where did it come from and how long has it been there? The answer to the first question is that the U.S. Department of Homeland Security believes it to be Russian made. The length of time it's been there is currently unknown but probably two to three years old. So, just because you have what you think is a good or great security architecture and you don't see any indications of a problem, doesn't necessarily mean that there isn't a threat lurking in your production network. You need a way to illuminate and remove the blind spots in your network so that you can understand exactly what is happening on your network and in every portion of it.

Finding The Perfect Solution: Combining Visibility and Security

This is where you will see some of the obvious benefits of blending the security and visibility architectures. In-line and out-of-band monitoring solutions can be created depending upon IT staff workload, and headcount, budget, etc. It primarily comes back

to what is the level of effort you want to invest. For instance, in-line monitoring solutions are extremely beneficial if you plan on taking a proactive security approach and want to deploy a SIEM, IPS, honeypot or so forth. If you don't have the staff or budget to support that level of effort, then you may want to deploy an out-of-band solution that takes a more reactive approach to your visibility architecture, but still gathers critical data to analyze potential security threats in a post-event, or time-delayed, situation. An out-of-band approach isn't real-time but it has the potential to be a lot less expensive and easier to deploy than an in-line approach.

Another consideration is to determine your level of risk. What do you have that you need to protect and how valuable is it? For instance, professional hackers aren't interested in attacking individual computers. What they really want is to hack routers, switches, intrusion prevention systems, firewalls, etc. The days of them just trying to attack the firewall are gone. At the same time, up to 70 percent of security-related events are due to insiders doing intentional or careless things that compromise the corporate network. So what is your highest threat level?

Proper network monitoring infrastructure and tools are essential to your network security. Without these capabilities it's sort of like looking at a house that has no windows or doors. Something is in there but what? A visibility architecture gives you the capabilities to create the door or window you need so that you can accurately capture anomalous information and mitigate the issue – whether it's a network malfunction and you need a fast mean time to repair (MTTR) or it's a security threat and you need to be able to visualize that threat.



The following capabilities can be deployed for production networks:


- ▶ Create in-line monitoring solutions for real-time delivery of anomalous traffic to your security solution
- ▶ Use out-of-band monitoring for routine analysis of network security
- ▶ Use web-based API's for implementing near real-time changes to your visibility network
- ▶ Gather application intelligence that provides DPI and additional value beyond Layers 2 to 4 packet data
- ▶ Audit for security policy infractions, including use of prohibited applications and devices
- ▶ Conduct troubleshooting and triage that employs iterative testing to help determine root cause

Implementing In-Line Visibility Solutions

For in-line solutions, these involve a tap that is inserted directly into the primary packet data path. You can then insert packet brokers and monitoring tools into the “loops” created by the in-line tap. The

data flows through these loops before being introduced back into the main data path for the network.

In the illustration below, an in-line (also called bypass) tap is placed in-line in the network for threat prevention, so it should be placed after the firewall but before any equipment. The advantage of this location is that should a threat make it past the firewall, that threat can be immediately diverted or stopped before it has a chance to compromise the network. Taps are also traditionally one of the few network devices that aren't susceptible to hacking since they usually don't have command interface capability.



A packet broker can be installed after the tap to provide filtering and load balancing capabilities, if needed. Important in-line tap and packet broker features include:

- ▶ Graceful load balancing fail-over mechanisms that do not disrupt existing sessions of available IPS appliances
- ▶ Load balance capabilities across a mix of different 1GbE and 10GbE appliances
- ▶ Automatic N+1 high availability load balancing

- ▶ Maintenance mode to allow convenient servicing of connected appliances in the load balanced group
- ▶ Ability to monitor multiple links in-line by load balancing across multiple IPS appliances

After the taps and packet brokers are introduced, in-line security tools are connected. These tools actively analyze the real-time data they are fed. Examples of typical in-line security tools include the following:

- ▶ Intrusion prevention systems
- ▶ Firewalls and next-generation firewalls
- ▶ Data loss prevention systems
- ▶ Unified threat management systems
- ▶ SSL decryptors
- ▶ Web application firewalls
- ▶ Deep capture devices used in data loss proof and historical studies

CASE STUDY #4

Financial Services & Bank (In-Line Monitoring)

BACKGROUND CUSTOMER is one of the largest international banking and financial services company in the world. They already had a DDoS solution and IPS solution but they wanted to increase their security defenses by adding a new IPS with increased capabilities, enhanced malware protection, and network forensics capabilities. The main drivers for the upgrades were revenue and reputation protection. Their three main goals can be summarized as follows:

- ▶ Add additional security tools in-line into their network
- ▶ Add high speed (40GbE) in-line taps and packet brokers with negligible network impacts
- ▶ Maintain ability to stop high-volume DDoS attacks

SETUP In-line taps were installed into the network. These taps serve as the primary access points into the network. Since they are placed in-line after the firewall and before the primary routing switches, they must have the ability to fail closed so that should they fail, traffic continues to flow through the network. The failover capability must also be extremely fast so as to minimize or prevent any loss of data. After the new tap, a specialized network packet broker was added to enable:

- ▶ Aggregation (from multiple taps)
- ▶ Load balancing (distribution of data to various security analysis tools)
- ▶ Filtering of data packets (for distribution to the correct security tool)
- ▶ Packet deduplication
- ▶ Packet slicing (delete non-relevant/sensitive information)

The taps and packet brokers also needed to handle high data rates (40GbE) and process functions like deduplication at line rate. In addition, CUSTOMER added an IPS (for IPS, DoS protection, and reputation services), a threat analytics system for malware protection, and a security analytics system.

RESULTS The solution worked as planned. After traditional DDoS prevention, traffic is now decrypted and forwarded to the IPS, and then sent on to the forensic recorder. Non-relevant information is deleted before it reaches the security tools to increase tool productivity and minimize network traffic. Fail-over capabilities were tested and performed as specified.

BENEFITS

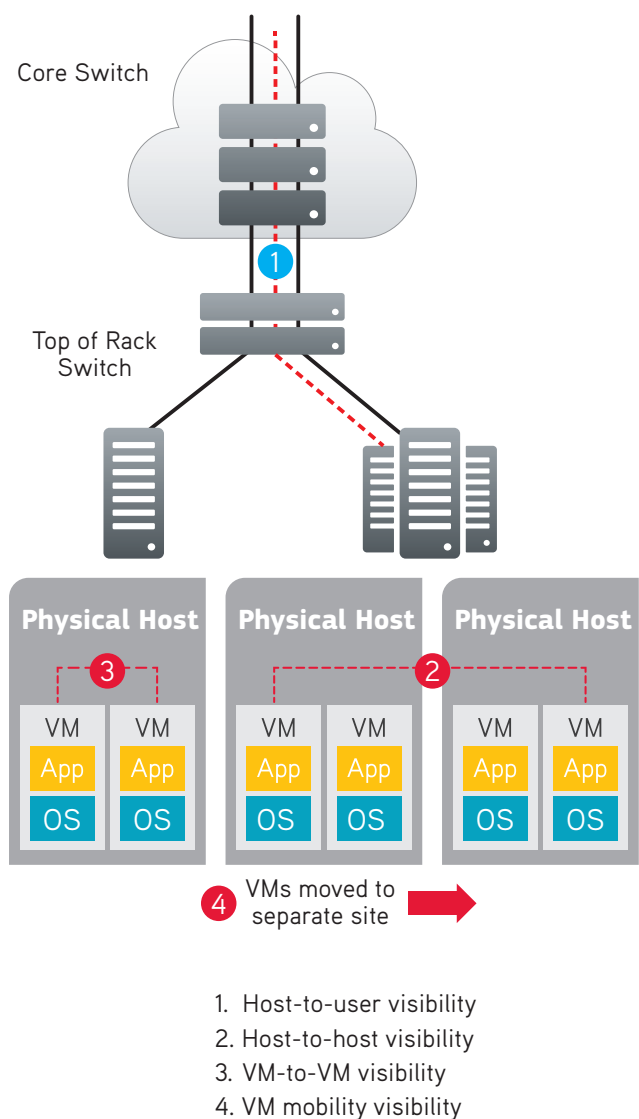
- ▶ CUSTOMER's network security architecture has been fortified
- ▶ The visibility and security architectures were seamlessly integrated to create a fast, in-line visibility solution that supports traffic at line rates
- ▶ The new capabilities successfully support existing capabilities without any network disruption

Extend Physical Network Monitoring Tools into the Virtual Realm with Virtual Taps

Virtual taps give IT monitoring personnel access to critical information within the virtual data center. The problem is essentially that while inter-VM (virtual machine) traffic was optimized to speed up connections and minimize network use on the physical core network switches, such optimization has made traffic invisible to physical tools, which are unable to extend easily into virtual environments.

In addition, next-generation data centers use virtualization technology to deploy private and public cloud environments on a single physical server or across a clustered group of servers, local and remote. Traditional taps cannot see the traffic between VMs that reside on the same hypervisor (east-west traffic), nor can they “follow” VMs as they are migrated from one host to another.

Visibility is further reduced by the complexity of blade servers that have each blade running multiple VMs on a hypervisor. Traffic running on blade servers shares a common backplane and creates a network blind spot, since the physical network and its attached tools are unable to see traffic above the virtual switch layer or the blade chassis network modules.



Virtual taps are specifically designed to enable protection for the virtualized data center. They are similar to standard taps in that they forward packets to network packet brokers and monitoring tools. The difference between the two types of taps is that the virtual tap is pure software.

CASE STUDY #5

Government Agency (Virtual Data Center Monitoring)

BACKGROUND CUSTOMER is a government agency for the United States of America. The main goal was to provide CUSTOMER with visibility into the virtualized data center to eliminate blind spots that could be hiding security and/or performance problems. They were particularly concerned because their existing monitoring solution only came into play for data that went beyond the top of rack (ToR). They were essentially blind to all inter- and intra-VM traffic and had malware concerns about the Crisis malware variant and other virtual environment threats. Their main goals can be summarized as follows:

- ▶ Insert virtual tap to gain visibility into traffic within the virtualized data center
- ▶ Export virtual data center traffic to packet brokers and monitoring tools for security and performance analysis
- ▶ Implement a common and consistent visibility solution across their whole network
- ▶ Access relevant virtual data center traffic to perform and demonstrate regulatory compliance (FISMA, HIPAA, GLBA, FEDRAMP, etc.)

SETUP Approximately 500 virtual taps were inserted into VMware 5.1 hypervisors. Filtering was turned on inside the virtual taps so that only relevant data was forwarded to the network packet brokers. This prevented overloading the LAN within the data center. The network packet brokers and monitoring tools were all existing. The virtual taps included an intuitive GUI interface which significantly reduced programming time and costs.

RESULTS The solution was implemented and virtual data center traffic began flowing to physical packet brokers. Once the monitoring data reached the packet broker, it was forwarded to the appropriate existing tools. For instance, a threat analytics system was already in place for malware threat detection and was now able to periodically monitor the virtual traffic (as well as the physical network packet data) for any threats. Data center traffic was also forwarded to a log file and event management appliance for regulatory compliance purposes.

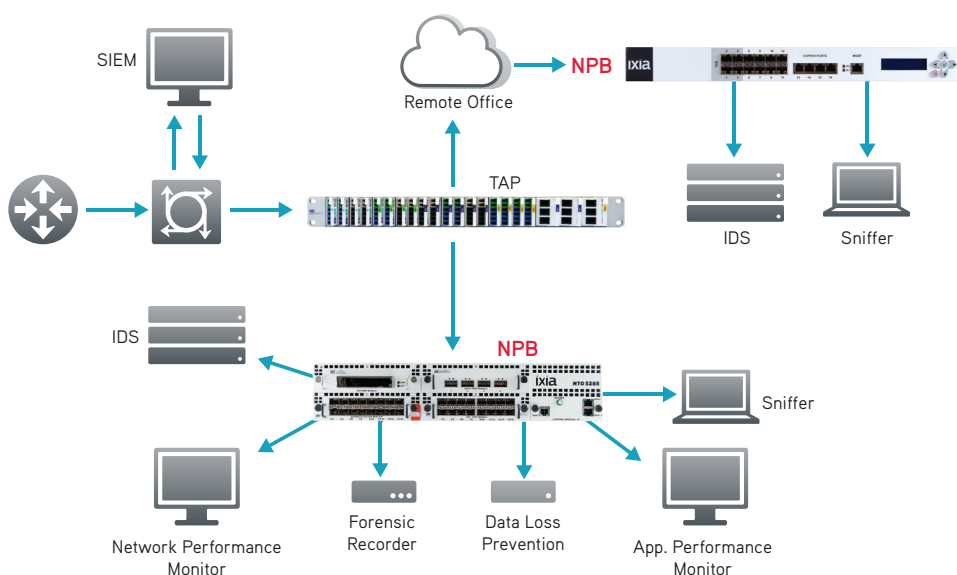
BENEFITS

- ▶ CUSTOMER now has visibility into all portions of their network
- ▶ Network security was strengthened as CUSTOMER is now able to periodically scan virtual data center traffic for malware
- ▶ CUSTOMER now has one consistent monitoring policy across their virtual and physical network segments that lets them demonstrate regulatory compliance
- ▶ CUSTOMER was able to leverage existing monitoring tools to control monitoring costs

Implementing Out-of-Band Visibility Solutions

As mentioned earlier, out-of-band solutions are less complicated and typically less expensive than the in-line solutions, but they still have distinct and necessary value. In fact, there is quite a bit of benefit here because the tools have more analysis time. For instance, forensic recorders can capture the initial information until other tools are ready to perform a full analysis on the threat/anomaly.

The following image shows how the NPB can distribute data to various tools in an out-of-band situation. Two packet brokers are shown in this example – one in the core and one at a remote office site.



Examples of typical out-of-band security tools include the following:

- ▶ Security information and event management (SIEM) systems
- ▶ Behavior analysis systems
- ▶ Forensic tools
- ▶ Data recorders
- ▶ Malware analysis tools
- ▶ Log management systems
- ▶ Packet capture tools

CASE STUDY #6

Healthcare Service Provider (Out-of-Band Monitoring)

BACKGROUND CUSTOMER is a children's hospital that had concerns over network security and the time it took to perform root cause analysis, which could go on for days. They also were interested in demonstrating HIPAA compliance and monitoring for their electronic medical record (EMR) systems which had documented slow response times. Their three main goals can be summarized as follows:

- ▶ Reduce root cause analysis time and improve performance trending initiatives
- ▶ Implement a better visibility architecture to improve tool access and remove duplicate data
- ▶ Improve the EMR system to increase patient satisfaction and comply with regulatory compliance initiatives

SETUP A network packet broker was installed after a combination of Cisco routers that already existed on the network. The packet broker was then connected to several monitoring tools including an end user quality of experience (QoE) tool. Deduplication and packet filtering was enabled on the packet broker using a drag and drop GUI interface.

RESULTS The solution was immediately used to help isolate the problem with EMR system – was the problem in the network or the application? Trending analysis from the packet broker/monitoring tool solution was able to isolate the problem for IT. The QoE monitoring tool also helped debug an email problem. The QoE trace data solved an Outlook/Exchange problem that had been plaguing the customer. General trouble shooting problem times also decreased, resulting in a MTTR of hours for most problems. This was an order of magnitude difference.

BENEFITS

- ▶ Troubleshooting time was reduced from days to hours
- ▶ Increased monitoring efficiency due to the eliminate of duplicate data and data filtering
- ▶ The EMR problems were eliminated, which also eliminated HIPAA concerns



Visibility Architecture Automation to Improve Security

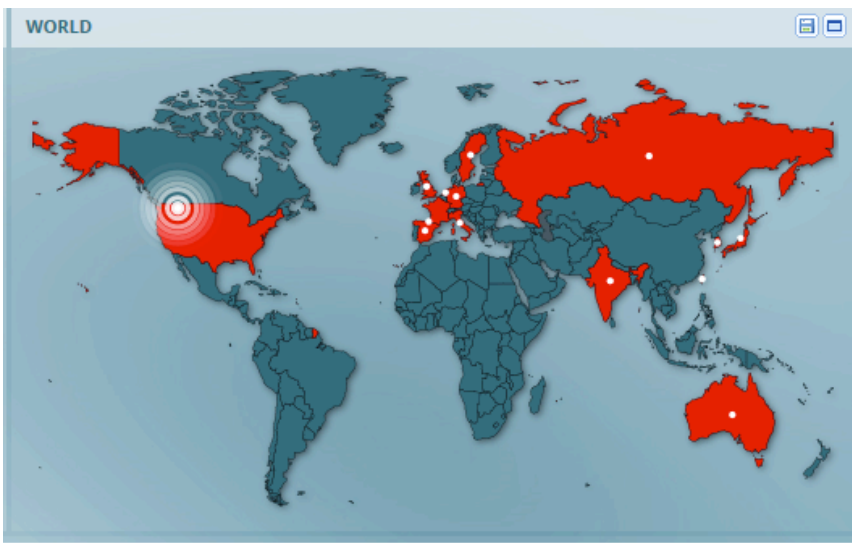
A third key feature of a combined monitoring and security architecture is to use automation capabilities. This is where the web-based APIs for packet brokers can be used to take full advantage of the packet broker's capabilities. This allows the packet broker to deliver near real-time capabilities in response to various stimuli, including anomalies and confirmed security attacks. Automating the network monitoring switch is one of the most powerful, but often overlooked, features for data center automation.

Enabling automation within your visibility architecture can deliver the following advantages to you:

- ▶ Real-time responses to mitigate or eliminate security anomalies and threats as they happen
- ▶ Faster responses to minimize the damage/cost to company

- ▶ Improved flow of information to and from intrusion detection and protection systems
- ▶ Improved flow of information to redirect threats to honeypots for better threat source isolation

The web-based API's allow packet brokers to communicate with network management and orchestration systems so that you can integrate an NPB with your data center automation initiatives. Near real-time changes in your visibility network (apply filters, add connections to more tools, etc.) can be created in response to external commands (to the NPB) to deliver data to your security tools. The source of the command could be a network management system (NMS), provisioning system, SIEM tool, or some other management tool on your network.

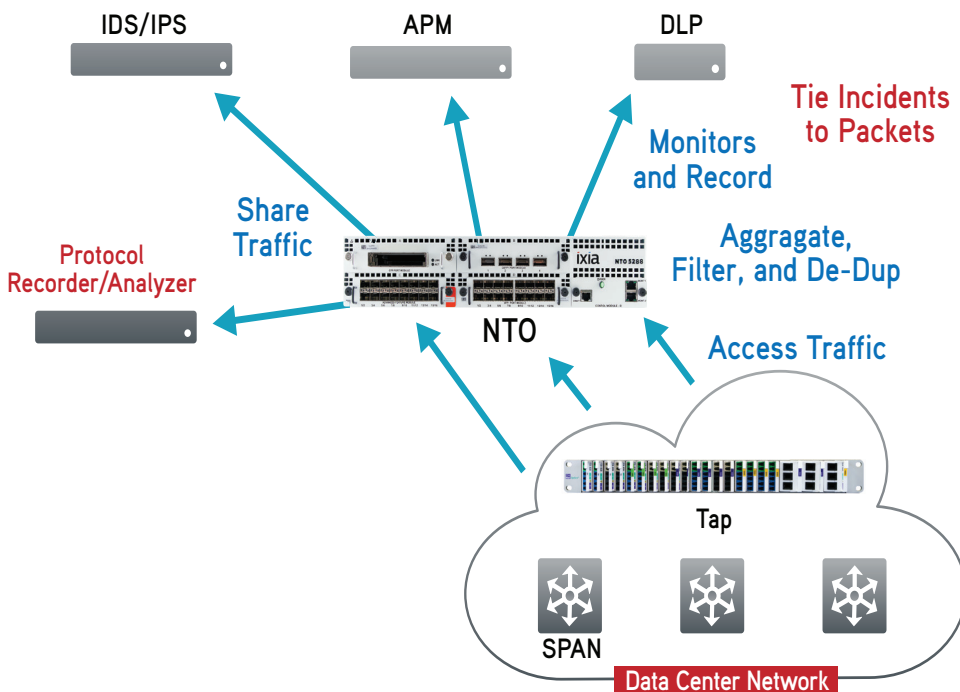


For example, should your SIEM detect that there is an incident taking place, it can direct the NPB to capture a certain set of packets on a certain link and then forward those packets to a specific security analysis tool for proper analysis. All of this would take place without any manual intervention with the NPB (assuming the NPB has the proper connections in place). It can also be used to dramatically decrease the MTTR for your network because faster responses to problems result in a shorter mean time to diagnosis and a corresponding faster MTTR.

Let's look at a quick example of how automation can help you. It's 3:00 am and you have a SIEM tool that spots some anomalous traffic on your network. The SIEM then sends a message to the monitoring switch instructing it to capture the anomalous traffic and send it to an IPS to analyze the anomaly. The packet captures are analyzed and this is identified as a threat by your IPS tool. Since you've previously connected your network monitoring switch in-line after your tap, the SIEM instructs the monitoring tool to divert this anomalous packet stream to a honeypot. From this honeypot you can now control the access that the intruder has and begin to understand better where the threat is coming from (e.g., who and where geographically), the threat attack used to gain entry into your system, and the nature of the attack the intruder had intended (defacement of the web site, crashing of the website, theft of corporate intellectual property, etc.).

In another example, the external equipment (e.g., ArcSight SIEM) sees an anomaly like a buffer overflow situation. The SIEM then sends a command to the packet broker to start a packet capture.

The filtered packet stream is then forwarded to another tool (e.g., an IDS tool) to determine whether this is a security attack, some random anomaly that may require further investigation, or some mistake (corrupted packets due to network or component failures).



If you have any further interest in this topic, there are actually quite a few use cases that have been detailed in the Ixia whitepapers [“Best Practices for Network Monitoring Switch Automation”](#) and [“Automation: The Future of Network Visibility”](#).

CASE STUDY #7

Global Energy Provider (Monitoring & Security API Integration)

BACKGROUND CUSTOMER is an energy provider and part of the national energy grid. They needed a system that could capture all anomalous behavior and quickly respond to security threats. They also wanted to accurately analyze and correct any issues. The existing manual diagnosis was slow and error prone. They had also maxed-out the connections to the monitoring tools and either needed to add an aggregator or additional (expensive) monitoring tools. Their four main goals can be summarized as follows:

- ▶ Quickly analyze anomalous behavior to determine any security threats and respond automatically
- ▶ Speed up time to resolution
- ▶ Reduce long term monitoring costs
- ▶ Simplify compliance reporting

SETUP The customer had multiple sources of network access – SPAN ports and taps. They added a network packet broker to aggregate the different inputs to their monitoring tools. The chosen packet broker also had a web-based API so that it could be connected to a new SIEM. The API interface was designed to allow the SIEM to issue commands to the packet broker, which could connect purpose-built filters to different network segments and then send the anomalous data to specific forensic tools for analysis. This would all be accomplished without manual intervention.

RESULTS The new architecture worked as designed. The automation capabilities were tested and they delivered the correct information to the monitoring tools in a fraction of the time it took for the previous manual process. This provided the customer confidence that the system could respond 24 x 7. The port aggregation and filtering also eliminated the need to buy any additional monitoring tools. The packet broker was also able to divert the appropriate data packets to their compliance reporting tools so that they could show due diligence for regulatory compliance.

BENEFITS

- ▶ An “always on” visibility solution was achieved through the integration of the security and visibility architectures to achieve the SIEM/NPB/forensic tool automation path
- ▶ Faster problem resolution (MTTR) was achieved
- ▶ Easier and faster access to monitoring and security tools at lower costs as no additional forensic tools were required
- ▶ Improved regulatory compliance reporting

NPB Application Intelligence Improves Security

Application intelligence is another advanced feature for network packet brokers that can deliver substantial value once the capability is integrated with your security architecture. Application intelligence is basically the real-time visualization of application-level data.

This includes the dynamic identification of known and unknown applications on the network, application traffic and bandwidth utilization, detailed breakdowns of applications in use by application type, and geo-locations of users and devices while accessing applications.

The filtered application information is typically sent on to third-party monitoring tools (Plixer, Splunk, etc.) as NetFlow information, but could also be consumed through a direct user interface in the NPB. The benefit to sending the information to third-party monitoring tools is that it often gives them more granular, detailed application data than they would have otherwise, improving their analysis. It also increases tool efficiency.

With application intelligence, you can reach all the way into Layer 7 (the application layer) of the packet data. It effectively allows you to create an early warning system for real-time vigilance. In the context of improving network security, application intelligence can provide the following benefits:

- ▶ Identify suspicious/unknown applications running on the network
- ▶ Identify suspicious behavior by correlating connections with geography and known bad sites

- ▶ Audit for security policy infractions, including use of prohibited applications

A core feature of application intelligence is the ability to quickly identify ALL applications on a network. This allows you to know exactly what is and is-not running on your network. The 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. Another key feature is that all applications are identified by a signature. If the application is unknown, a signature can be developed to record its existence. These unknown application signatures should be the first step as part of IT threat detection procedures so that you can identify any hidden or unknown network applications and user types.

A second feature of application intelligence is the ability to visualize the application traffic on a world map for a quick view of traffic sources and destinations. This allows you to isolate specific application activity by granular geography (country, region, and even neighborhood). User information can then be correlated with this information to further identify and locate rogue traffic.

For instance, maybe there is a user in North Korea that is hitting an FTP server in Dallas, Texas and transferring files off network. If you have no authorized users in North Korea, this should be treated as highly suspicious. At this point, you can then implement your standard security protocols, for example kill the application session immediately, capture origin and destination information, or capture file transfer information.

You can also use application intelligence to audit your network policies and ongoing usage of those policies. For instance, maybe your official policy is for employees to use Outlook for email. All inbound email traffic is then passed through an anti-viral/malware scanner before any attachments are allowed entry into the network. With application intelligence, you would be able to tell if users are following this policy or whether some are using Google mail and downloading attachments directly through that service, thus bypassing your malware scanner. Not only would this be a violation of your policies, it presents a very real threat vector for malware to enter your network and commence its dirty work.

CASE STUDY #8

Wireless Services Provider (Application Intelligence)

BACKGROUND CUSTOMER is one of the top 3 international wireless service providers. The global service provider needed the ability to understand exactly which applications were running on their network and any associated security or performance issues with specific applications. Application intelligence was desired so that CUSTOMER could get rich data on the behavior and location of users and applications. The data needed to be created and exported in any format required by the monitoring analysis tools – raw packets, filtered packets, or NetFlow information. CUSTOMER had experienced a network outage due to massive, rapid bandwidth consumption by a user application and wanted to make sure that did not happen again. Their three main goals can be summarized as follows:

- ▶ Improve network security by identifying hidden applications running on the network
- ▶ Correlate user and device geolocation with specific applications to improve network security
- ▶ Reduce network outages and/or improve network performance with application trend information

SETUP An application and threat intelligence processor was installed into a network packet broker on CUSTOMER's network. It was connected into a tap after the core routing switches. NetFlow output was generated by the application intelligence processor and fed to Splunk, which provided a visual display of the applications in use, geolocation of applications and devices, and trending information for application usage.

RESULTS T teams were able to identify three previously unknown applications on the internal network. On the production network, applications were recognized and data fed to Splunk so that performance and trending charts could be created for a weekly review to make sure that there are no more network outages.

BENEFITS

- ▶ CUSTOMER is now able to quickly identify suspicious applications and devices and then correlate with connections to known bad sites
- ▶ CUSTOMER is able to identify all applications (known or unknown) that are running on their customer network and proactively manage that bandwidth
- ▶ IT was able to implement better network security to find unauthorized applications
- ▶ CUSTOMER can conduct random audits of application usage and devices on their corporate network to see if there may be any policy violations happening

Chapter 6:

Software and Hardware Upgrade Testing Limits Security Flaw Introduction

So you've gone through design and testing and your network is now up and running smoothly. But we all know that this is not the end of the process for vibrant, functioning, and secure networks. Whether due to business needs to expand network functionality, or external influences like vendor software updates or the discovery of new malware or vulnerabilities, your network infrastructure will change – and it will change often.

This final stage of the network lifecycle is the change management phase. This is also a very dangerous stage as every change in the network is a potential threat to your existing network design that has already been tested and validated! As network hardware and software upgrades are implemented, you need to ensure an appropriate security posture throughout the changeover. An essential discipline for IT, change management involves a formal process to request, approve, implement, and audit network changes, balancing the rapid delivery of software changes while ensuring that performance and security remain intact.

Change management is, in effect, threat management to uncover weaknesses in the existing devices and software, but it also presents new risks and unknown parameters for security issues and attacks. Each change, even those required to overcome vulnerabilities, has the potential to introduce its own new vulnerability.

Every change should be fully vetted and tested against all applicable hardware, software, and firmware. The upgrade process is one of the most dangerous to handle as it needs to be verified and

tested throughout the process and the systems affected need to be revalidated and re-commissioned, just like starting over.

Pre-testing of upgrades is instrumental in preventing the introduction of new network security flaws. As new patches, updates, revisions, etc. become available, validate them and compare the results to your most recent baseline data. Empirical data from the testing will allow you to make better-informed decisions about what you should put into production.

Upgrades – Where Dangers Lurk

While some might think this stage sounds like a lot of unnecessary work, it's actually much-warranted as this is probably the most dangerous phase of your network's lifecycle. Upgrades, especially software updates, are one of the most successful attack vectors against a network. They are also fairly common and affect a lot of users and systems. Java, Flash, and Acrobat reader are just a few of the common upgrade attack points. The Heartbleed vulnerability, for instance, is a recent example of this attack vector. It was introduced as part of a software upgrade for SSL, but it actually exploits the heartbeat extension. A follow-on to Heartbleed is the Cupid attack that works in a similar manner to affect Wi-Fi TLS services, particularly for WPA clients and servers.

Here's a sampling of just a few upgrade attacks that were found in the last couple years and mentioned in a [Dell paper](#):

- ▶ IBM Tivoli Provisioning Manager Express SQL Injection
- ▶ VideoLAN VLC Media Player mms Buffer Overflow

- ▶ Wells Fargo Account Update Downloader Trojan
- ▶ New LockScreen Ransomware Trojan in the wild
- ▶ Oracle Java Runtime TTF BO

Not only are upgrade attacks irritating, they can become expensive. Besides the typical costs for personnel to remediate the malware and it's after affects, there are other subliminal upgrade attack types, like malicious video posts, which inform you that your video player codec is out of date and ask you to upgrade your player software. Once you agree, malware, called "ransomware" is downloaded to your computer or server which encrypts your hard drive or locks up your computer until you agree to pay the malware creator a ransom. Once you accept the upgrade request, the malware loads in the background and then executes itself to encrypt your data. You are then prompted to pay a ransom to get your own data back. If you don't pay, you lose everything or you have to pay someone else a lot of money to try to recover the data, which could easily be a futile effort. If you decide to pay, the attacker doesn't always send you the key to unencrypt your system, so you're out the money and still have a worthless computer. In the case of ransomware, you probably won't be encountering this much with your production servers but it's possible that a network user could introduce the threat into your production network.

3 Steps to Reduce Change Management Risks

Here are 3 steps we highly recommend to help you mitigate change management risks:

Step 1: Validate Upgrade Authenticity

There are two points to make here. One is about making sure your change has been authorized and the other is making sure you're implementing untainted software.

An effective change control process comes into play here, as it helps you determine which change requests have been authorized and which are not authorized. Opening a TCP port on a firewall to help a traveling employee who can't access a host can cause a security breach that results in a catastrophic attack. Strong suggestion – do not implement changes that have not gone through your formal change process.

Before making any change on your system, it is also strongly suggested that the network manager evaluate the software upgrade for authenticity. Supply chain integrity minimizes the risk of malicious code. Make sure to use your application's interface or the vendor's specific download site for updates and upgrades.

Additionally, you'll want to verify the SHA (secure hash algorithm) and memory read-off of the processes to determine if it matches specifications from the manufacturer. You can use the **file verify auto** global configuration command to verify the integrity of

some software images. Or manually, simply record the hash as presented by the vendor's upgrade tool and then once the download is complete, the hash of the local file should be verified against the hash you recorded. Once the software has been verified as authentic and unaltered, copy it to write-once media and verify the hash of the copied file to detect corruption during the copy process. Remove the local file and move the read-only file to the file server, verifying the hash once again. A few extra minutes here can save hours of time and embarrassment later on.

Step 2: Validate in the Lab – Interoperability and Security

Once you're ready, establish what the changes and effects will be on your operations and users in a lab setting before upgrading the production system. This evaluation includes how the changes will affect the operating system of the device as well how it will affect other approved software and hardware products currently in use.

As part of your change verification process, you'll want to definitely consider these areas:

- ▶ Establish effects of changes on the network, users, and any vulnerabilities
- ▶ Establish other network, security, or user dependencies caused by the changes
- ▶ Establish the different domains and hierarchy for changes

This step requires the use of a comprehensive security assurance and validation solution that includes the ability to:

- ▶ Simulate past and recent attacks

- ▶ Generate real-world, line-rate application loads
- ▶ Assess devices and entire systems
- ▶ Scale to cover extensive test cases

Since new vulnerabilities are discovered on a daily basis and software is continually evolving, it is critical that your security assurance solution is always current and ready for action; so keep it up-to-date with the latest software releases and security attacks.

When actually testing the upgrades, you'll probably find it easier if you have implemented a high availability (HA) visibility architecture. The HA architecture provides redundant in-line taps, packet brokers, and other equipment so that you can validate the updates for various portions without taking all of your key security equipment out of service at one time. This is a best practice that can be combined with your out-of-band visibility architecture to perform upgrade validations as quickly and easily as possible.

Step 3: Validate Devices and the Full Network

After you've implemented your changes, make sure you test any changed devices and the fully commissioned network. This is usually done by retesting against the last baselines to establish change effects. Once the changes are made and accepted, save the new baselines for alarms and the comparative parameters for your network management system.

Baselines should include:

Metric	Key Performance Indicator (KPI)
Performance	<ul style="list-style-type: none">> Connections/sec> Total connections> Number of simulated users> Throughput
Application-Level Transactions and Failure Monitoring	<ul style="list-style-type: none">> Requests sent/successful/failed> Request aborted> Timeouts> Session timeouts> Connect time
TCP Connection Information	<ul style="list-style-type: none">> SYNs sent> SYN/SYN-ACKs received
TCP Failure Monitoring	<ul style="list-style-type: none">> RESET sent> RESET received> Retries> Timeouts
DoS Attacks	<ul style="list-style-type: none">> Successful, failed packets> Bytes sent
Packet Monitoring	<ul style="list-style-type: none">> Packets received> Packets filtered> Packets allowed

If you are like most IT/security professionals, you’re interested in anything that can expedite in-house network security testing. One innovative option is to use “test packs” from third-party test labs that include the methodology for testing security devices like IPSs and firewalls. Using test packs provides a shortcut to benchmarks and information regarding security effectiveness, performance, manageability, and cost of ownership.

Best Practices for Secure Change Management

Here are some additional recommendations to help prevent the introduction of security flaws into your architecture:

- ▶ Establish and document a change management plan (when are you going to upgrade, how often, what verification and installation protocols will you follow, etc.). Organizations that do not have a change management process face poor software quality, failed changes, rework, network downtime, dissatisfied customers, missed deadlines, and higher costs.
- ▶ Establish and correlate levels of change to your risk level and then secure approval from management
- ▶ Establish a formal method of change notification to users and how you will distribute the actual change
- ▶ Establish change installation management from process and procedure
- ▶ Establish change verification and test process
- ▶ Keep a record of who made changes and when they were made
- ▶ Compare your processes to international standards like [ISO 27001](#), [27002](#), [27005](#) to make sure they are best practices for IT governance and infrastructure security and network compliance. More information on the ISO 27001 processes can be found [here](#).

CASE STUDY #9

Financial Services Provider (Ongoing Device Testing)

BACKGROUND CUSTOMER is a large financial enterprise that has deployed a NGFW. Much thought and analysis went into making the decision on which NGFW to purchase. Additionally, CUSTOMER wanted ongoing device testing to ensure updates didn't impact performance. When the NGFW project was deployed the goals were:

- ▶ Improve application inspection and policy enforcement
- ▶ Deploy devices that could perform against their goals for application inspection
- ▶ Provide better protection and prevention for current and evolving threats

SETUP CUSTOMER did a proof of concept (PoC) and full evaluation of multiple vendors before making an NGFW product decision. The product was selected that met their performance goals and had very effective detection and prevention of threats. Testing allowed them to validate their expectations and the capabilities communicated by the vendor.

RESULTS CUSTOMER was able to deploy the NGFW product after their testing was completed. After it was deployed, months later an update was provided from the vendor. New functions and features were noted but nothing to indicate the product effectiveness had changed. Upon testing, it was discovered that the security effectiveness had changed. The NGFW was now missing attacks and attempts to evade detection that had been tested earlier were being missed.

CUSTOMER realized that something had changed and upon more research realized that indeed some third-party test reports had seen similar issues with vendor product versions and security effectiveness changes.

Ongoing testing before introducing new updates and versions of software for their security device has allowed them to verify its effectiveness and ensure that an update does not introduce new vulnerabilities or holes in their security strategy.

BENEFITS

- ▶ Ongoing testing identified issues introduced in updates
- ▶ Holes or gaps in security effectiveness were prevented
- ▶ Product vendor was advised and new updates made available to CUSTOMER

Chapter 7: Conclusion

IT and business decision makers today face various threats to the security of the enterprise. Approximately 43% of companies are experiencing data breaches on an annual basis (according to [USA Today](#)) as an onslaught of corporate security attacks is accelerating. A Ponemon study ([Ponemon Institute's 2014 Cost of Data Breach Study: Global Analysis](#)) also indicated that the average cost to a company for a breach was approximately \$3.5 million in US dollars, which is 15% more than what it cost last year.

Besides the acceleration of security attacks, other bad news exists as well. Of the security breaches reported, 66% of the breaches were discovered more than a month after they occurred (according to the [Verizon 2013 DBIR](#)). What's more distressing is that 71% of the breaches weren't even detected by the company that was breached – someone else had to report to them (according to the [2014 Trustwave Global Security Report](#)). This doesn't bode well for enterprises that rely on traditional network security architectures and processes.

A systems view that uses a network lifecycle approach offers an alternative to the common pitfalls in today's security architectures. The network lifecycle approach offers the following benefits at each of the four stages:

- ▶ Pre-deployment security analysis validates your security design
- ▶ Threat assessments during the installation and commissioning stage ensures design compliance
- ▶ Production network visibility enables real-time security vigilance

- ▶ Software and hardware upgrade testing limits, if not prevents, security flaw introductions

Using the network lifecycle approach also allows you to clarify your goals and processes. Specifically, it will help you:

- ▶ Define your specific vulnerabilities (servers, network devices, services, access policies, etc.)
- ▶ Establish loss risk levels
- ▶ Define your idea of what an attack in your network looks like
- ▶ Learn and understand the most common attack vectors that could be used against you
- ▶ Validate that your visibility/recognition plan is tuned towards your most likely attack
- ▶ Define a mitigation plan for most attacks types
- ▶ Create a remediation plan
- ▶ Develop a proof of loss plan to prevent large civil penalties
- ▶ Create a retest, revalidate and commissioning plan

Here’s a simple table to consolidate suggested actions.

Stage	Key Actions
Pre-Deployment	<ul style="list-style-type: none">› Define your specific vulnerabilities (servers, network devices, services, access policies, etc.)› Establish loss risk levels› Define your idea of what an attack in your network looks like and a mitigation plan› Create visibility architecture› Create security architecture› Validate proposed architecture components to make sure they work as specified› Validate your proposed architecture using test equipment› Learn of and understand the most common attack vectors that could be used against you
Installation and Commissioning	<ul style="list-style-type: none">› Create a remediation plan› Implement network changes› Validate that your visibility/recognition plan is tuned towards the most likely attacks› Develop a proof of loss plan to prevent large civil penalties
Production	<ul style="list-style-type: none">› Implement remediation plans for security attacks and breaches› Run periodic checks on equipment logs to investigate for breaches› Perform periodic scans for anomalies› Perform trend analysis and compare periodically to the baseline› Notify senior management of breaches according to published procedures
Software and Hardware Upgrades	<ul style="list-style-type: none">› Create a retest/revalidate and commissioning plan› Validate software and hardware upgrades using test equipment› Update processes, procedures, and your baseline based upon changes to your architecture

In the end, the key point should involve a mindset change towards the philosophy that network security is something you will need to do on a daily basis, not just at a point in time.

The background features a dark, swirling pattern with vertical columns of binary code (0s and 1s) in a lighter gray. The 'ixia' logo is prominently displayed in the center-left.

ixia

26601 W. Agoura Road
Calabasas, CA 91302 USA
Tel +1.818.871.1800
www.ixiacom.com

January 2015 – 915-3527-01 Rev. A