



ANALYST BRIEF

Cyber Resilience

IT'S NOT ABOUT THE 98 PERCENT YOU CATCH, IT'S ABOUT THE 2 PERCENT YOU MISS.

Authors – Bob Walder, Chris Morales

Overview

Where the goal of cyberprevention has been to reduce the probability of an attack against the organization, cyber resilience looks to reduce the impact of these attacks through cyber risk management. A cyber resilience program still considers detection and prevention techniques, but it also assumes that a breach is likely. This stance emphasizes anticipation, agility, and adaptation. Not every attack can be prevented, but with a cyber resilience program, damage can be mitigated or avoided altogether.

NSS Labs' research has determined that most modern leading security products can generally be considered to be effective. For example, during the 2013 network intrusion prevention system (IPS) group test, average security effectiveness (which factors in exploit block rate, anti-evasion capabilities, and stability/reliability) across the ten products tested was 94 percent. During the 2013 next generation firewall (NGFW) group test, eight out of nine products scored more than 90 percent for security effectiveness. The highest security effectiveness score in that test was 98.5 percent.

However, it is not the 98.5 percent that is caught that is the issue; it is the 1.5 percent that is missed. If even a small fraction of that same 1.5 percent of current threats is missed by the NGFW, IPS, and endpoint protection (EPP) system, then we have the beginnings of a breach.

Current cyberattack campaigns involve stealthy, persistent, and sophisticated activities to establish a foothold in organizational systems; maintain that foothold and extend the set of resources the adversary controls; and exfiltrate sensitive information or disrupt operations.

Enterprise architecture and systems engineering must be based on cyber risk management principles in order to ensure that mission and business functions will continue to operate in the presence of compromise.

NSS Labs Findings

- It's not the 98 percent that is caught that matters, it's the 2 percent that is missed.
- Cyber resiliency of systems and networks is needed to ensure mission survivability in a cyber-compromised environment.
- Security controls should be viewed not as complete protection against attack, but rather as a means of maneuvering the adversary into attacking a target of the organization's choosing, and also as a means of proactively managing the impact of network penetrations.
- It is not sufficient to perform vulnerability scanning to map potential risks. Actual testing of security devices is required to provide empirical evidence of their weaknesses.

NSS Labs Recommendations

- Determine the answers to critical questions such as:
 - What are the attacks that are being used by threat actors today?
 - Which of those attacks are effective against the business applications deployed in the network?
 - Which of those attacks are capable of bypassing deployed defenses?
- Focus on key metrics such as *time to awareness* and *time to response*.
- Learn to anticipate attacks. Assume the breach will occur, and focus on reducing its potential impact.
- Prepare to operate at 60 percent capacity in order to withstand a breach, which will reduce, but not eliminate, critical services.
- Plan for flexible network architectures that will allow dynamic reprovisioning of critical resources to isolate and replace infected portions of the network.
- Segment networks so that a compromised low-priority host cannot infect the rest of the organization and cause system-wide loss.
- Do not seek to remediate a breach immediately. Isolate the infected portion of the network and learn why the attack was successful while it is still underway, then redesign the architecture to withstand similar attacks.
- Increase the cost to the threat actor and buy time for investigation and remediation through the use of deception technology.

Table of Contents

Overview	1
NSS Labs Findings.....	2
NSS Labs Recommendations	2
Analysis	4
Current Best Practice May Not Be Enough.....	4
Problems with Traditional Vulnerability Scanning and Threat Intelligence	5
Moving from Reactive to Proactive.....	6
From Secure Networks to Resilient Networks.....	7
Resilience as the End State.....	8
Reading List	10
Contact Information.....	11

Analysis

Best practice for many years has been to deploy multiple types of security products both at the network level and on endpoints in a “defense-in-depth” strategy. This strategy has been based on two important assumptions. The most obvious is that there was a clear distinction between “inside” and “outside” the corporate network. Building high walls and digging a moat only makes sense if the crown jewels are locked in the keep.

Another assumption is that layering security protection will significantly enhance overall security. Mathematically, this is represented by calculating the overall failure rate of a combination of security products, which is achieved by multiplying the failure rates of all of the products. For example, if an organization’s IPS allowed one attack in 100, and its NGFW allowed one attack in 100, then the combination of these products should allow only one attack in 10,000 (0.01×0.01).

However, NSS research has determined that this calculation significantly overestimates the level of protection achieved by defense in depth, because there is strong correlation in the threats that pass through existing classes of products.¹ Further, the rise of cloud services and the widespread use of mobile devices have made the use of traditional defense-in-depth architectures far less attractive because data now regularly migrates outside the traditional walls of an organization.

Current Best Practice May Not Be Enough

Current best practice is to promote defense in depth with preventative security controls, but also to add breach detection and remediation components with the expectation that breaches eventually will occur.

Practitioners often describe security as a process that involves managing the appropriate actions before, during, and after a breach. This process includes designing, monitoring, and managing appropriate security controls; continuously monitoring internal network traffic to quickly identify a breach; remediating any breach damage; and modifying security controls based on new information. The process then repeats.

Focus no longer should be on the 98.5 percent of attacks that security controls currently are detecting and blocking, but rather it should be on the 1.5 percent of attacks that pass unnoticed through existing defenses, and on whether those missed attacks are relevant to the organization. Note that all security technologies, even breach detection technologies, fail to detect a significant number of attacks. Knowing which types of attacks they miss (i.e., Mac OSX or 64-bit Windows) is key to managing risk.

The extensive ecosystem of security controls deployed today should not be viewed as complete protection against attack. Rather, it should be viewed as a means of maneuvering the adversary into attacking a target of choice and then proactively managing the impact, as well as reducing actual network penetrations to a level that can be managed by remediation and response teams.

Enterprises constantly must assess the trade-offs between increased productivity and cost savings associated with the use of public networks, and also the security concerns that arise from this business decision.

¹ <https://www.nsslabs.com/reports/modeling-exploit-evasions-layered-security>

Securing the nation’s most critical infrastructure will require more drastic measures, such as providing an air-gap, or combination of unidirectional security gateways and secure bypass devices, between critical infrastructure and public networks in order to reduce exposure to threats.

Enterprises should also investigate the use of deception technology to increase the cost to the threat actor (possibly deterring or foiling less committed threat actors) and buy time for forensic investigation and remediation.

Problems with Traditional Vulnerability Scanning and Threat Intelligence

While traditional threat intelligence offerings, threat feeds, and vulnerability scanning practices remain useful tools in the struggle against cybercrime, they have significant flaws that make them of limited use to enterprise security teams.

Threat feeds will inform when a new vulnerability is discovered (less useful), or when a new exploit is seen in the wild (more useful), or sometimes both. In other words, they can alert the security team to a potential threat, but they cannot pinpoint its target. The risk is that organizations will rush to update to the latest version of the application for which an exploit has been reported, only to discover that they are now more vulnerable than before.

For example, previous NSS EPP group tests showed that while the protection for Java 6 update 23 was nearly 100 percent, it was less than 5 percent for Java 7 Update 2. A patch management policy that forces an update within a limited time frame without regard for security product capabilities could leave systems more exposed than they were before the update. Smarter patch management that takes security product capabilities into account is critical.

Vulnerability scans, on the other hand, are a snapshot in time and will reveal all of the possible vulnerabilities in the applications deployed within an organization. What they are not able to discern is how many of those potential vulnerabilities are effectively neutralized by the security defenses currently deployed. This can generate a laundry list of problems for the security team, most of which could be a much lower priority if the full facts were known.

Focusing on log entries and security information and event management (SIEM) data can result in a backward-looking stance that will not help defend organizations against the more sophisticated threat actors. The key question is not “*what was caught?*” but “*what got through?*” The problem is that security products do not provide logs and reports on the attacks they miss. So trying to determine where a breach might have occurred and what damage it might have done is akin to searching for the proverbial needle in the haystack. Except in this case the needle is not even in the haystack.

In an organization with security controls, threat actors will look for weaknesses in its security products. From this perspective, security products are “pushing” attackers away from targets where security is strong and toward others where security is weak. Compromises occur when attackers target an application used by an organization, and security controls fail to stop the attack. In this case, it is imperative for organizations to understand which attacks are not blocked by the security products they have deployed if they are to minimize exposure and reduce risk.

In order to predict where threat actors will focus their attacks in advance of a breach occurring, security professionals will need to determine several critical pieces of information:

- **Accurate details of the attack methods and exploits being used by cybercriminals.** This should include the types of assets being targeted (server or endpoint, specific business applications, and so on) in order to determine if the organization's attack surface is vulnerable.
- **Details of information technology (IT) assets and their value to the business.** This data can then be matched against the attack methods currently in use by threat actors to determine if IT assets are indeed at risk (for example, an endpoint with no instance of Adobe Air is not at risk from an Adobe Air exploit).
- **Effectiveness of the security controls deployed.** If an organization is at risk from a particular Java exploit but the IPS deployed is capable of identifying and blocking that exploit, the organization is not at immediate risk.
- **Likelihood of the organization being targeted directly by threat actors.** Financial institutions, defense organizations, and critical infrastructure installations have a higher risk of being targeted by threat actors than small manufacturing concerns.

Following any breach, the key metric is *time to awareness*. However, it is not enough to be aware only that malware was dropped on a host, since this is often easily remediated and that remediation can provide a false sense of security if the breach is more extensive. It is more important to be aware of the moment intellectual property (IP) leaves the network perimeter, or the moment a business-critical service fails. Security teams that focus purely on blocking malicious traffic are missing the real issue. It is not the initial attack that is important, nor even the subsequent infection. It is the interruption in service or loss of IP that is crucial.

Organizations should assume the breach will occur and proactively seek to reduce the impact of that breach. That is the key to cyber resilience. True cyber resilience allows organizations and governments to continue to operate and provide services for clients or citizens in the face of persistent and never-ending attack. Instead of trying to stop attacks in cyberspace or even at the network perimeter, networks must become resilient so they continue to function regardless of the level of attack.

Moving from Reactive to Proactive

Existing compliance requirements are often criticized by security practitioners as static, brittle, and disruptive to other security priorities. The Payment Card Industry Data Security Standard (PCI DSS) is a good example. The PCI DSS establishes baseline requirements for businesses that store, process, or manage credit cards. PCI DSS audits provide point-in-time assessments of a merchant's security posture, but what is needed is continuous monitoring and risk assessment.

Ideally, vulnerability scan audits of security controls would move from point-in-time to continuous, creating a real-time model. Currently, however, there is no business or industry set up in this way. In addition, there is a fundamental flaw in many existing methodologies: they do not take into account the capabilities of the threat landscape as it exists at any given point in time (are exploits in the wild targeting the attack surface deployed within the organization?) Nor do they take into account the capabilities of the security devices currently deployed.

Only by carefully examining the exploits currently being used to deliver malware to an endpoint can their effectiveness be determined against any particular endpoint deployment model. And only by combining that knowledge with the security effectiveness of the protection products in place (IPS or EPP, or both, for example) can it be determined if the endpoint is actually at risk.

This knowledge can be used to calculate a risk level that is based on business profile (including the likelihood that the business will be targeted by threat actors), services dependence, asset criticality, and operational needs.

For example, a current exploit may be effective against Internet Explorer 9 on Windows 7 Service Pack 2 (SP2) but not against IE9 on Windows 7 SP1. If an enterprise knows its current attack surface (the applications deployed on its endpoints and servers), then the **potential** risk level can quickly and accurately be determined. Further, if the enterprise also knows exactly which exploits bypass its IPS and EPP products, the **actual** risk level can be calculated.

If the IPS/EPP combination blocks the exploit, the attack surface is irrelevant. The organization has a potential problem, but it can patch at leisure while being protected by its security mechanisms. If the IPS/EPP combination does not block the exploit, but the organization is not running the affected application (i.e., it is running SP1 instead of SP2), then it is not in immediate danger of being breached. However, it now has early notice of the need to update its security products before updating to SP2.

On the other hand, if the IPS/EPP combination does not block the exploit and the organization is already running SP2, then the risk is extremely high, and the risk level will reflect this. Any attacker using that exploit against that particular endpoint configuration will be successful, and a breach is likely to occur. Armed with this information, enterprise security teams can address security issues proactively, moving swiftly to ensure that vulnerable products are updated or replaced, and also to verify that the system is clean (i.e., that no breach occurred).

In this case, log files would be of limited value, since if a security product is vulnerable to a particular exploit, how will it ever be in a situation where it can log anything? This is a situation where knowledge of a problem (for example, an exploit is in the wild and currently being used by threat actors; the exploit is effective against a particular endpoint configuration in the organization; and that same exploit is capable of bypassing all of the currently deployed perimeter and endpoint defenses) can provide enough notice to enable a security team to take preventative measures in advance of a breach.

Knowing that the network was vulnerable to such an exploit even for a short time would also allow the security team to focus its forensic investigation efforts in the right areas to determine whether a breach actually occurred.

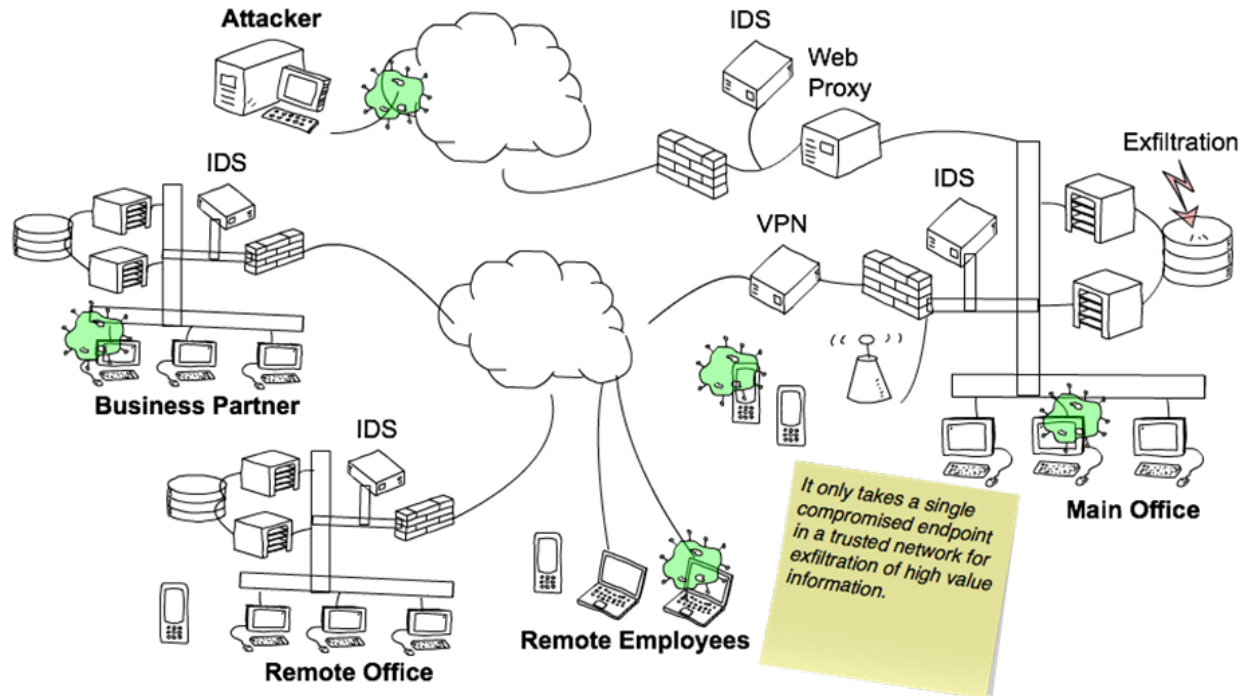
From Secure Networks to Resilient Networks

The manner in which enterprise information is used, accessed, and stored is changing with virtualization, cloud computing, mobility, consumerization of technology, the “Internet of Things,” and the deployment of critical production systems. Enterprises must adapt to these changes in order to remain competitive. As new technologies are adopted, such as software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS), they introduce new risks, and information security departments must evolve to protect existing on-premises, legacy systems against them.

NSS research reveals that the traditional enterprise information security architecture (EISA), which uses network perimeter and asset-based security models, can no longer provide the controls and countermeasures that an organization requires in order to address these risks. Trusted computing models must be redefined to limit the exposure of enterprise information and protect critical infrastructure. Architectures must be based on the assumption that systems will be compromised, but businesses must continue to function.

Cyber resilience provides the framework for a new architecture that is appropriate for these new environments. Where the goal of cyberprevention has been to reduce the probability of an attack against the organization, cyber resilience looks to reduce the impact of these attacks and provides the ability to operate in the face of persistent attacks. Resilience enables the government to continue to provide services to the public, and industry to continue to serve employees and customers while repelling cyberattacks.

A cyber resilience program still considers detection and prevention techniques, but it also assumes that a breach is likely. This stance emphasizes anticipation, agility, and adaptation. Not every attack can be prevented. But with a successful cyber resilience program, damage can be mitigated or avoided altogether.



Resilience as the End State

The modern world depends on services that are expected to be available 24x7x365; for example, mobile phone networks, electricity control mechanisms, media providers, and large online shopping sites. The security controls that protect these assets should also be available 24x7x365, but initially, many were static, which meant that security holes could be exploited for months or even years. Today, many security technologies possess dynamic control mechanisms with near real-time signature and heuristics updating, IP blacklist management, and isolation and remediation mechanisms.

To achieve cyber resilience, tradeoffs are often made. Each additional level of security impacts the user experience, or business performance. Further, some assets may not have the ability to update while in use and will require time offline, which means they will not provide continuous protection.

Dynamic provisioning is one approach to resilience, since it offers the ability to continue providing services on an infected network. This requires a network smart enough to reprioritize traffic and to rearchitect itself on the fly, isolating the infected portion of the network into a contained area. At the same time, new resources can be added to the network to reroute traffic and manage it outside the infected environment.

This approach is increasingly viable with the adoption of elastic cloud architectures. The aim is to reduce the *time to response*: once the infected portion of the network has been contained, services can continue as before, and the security team can analyze the attack (which can continue within the isolated environment) and determine the most appropriate remediation steps.

The key goals of this approach are:

- **Anticipate.** Assume the breach will occur, and focus on preparedness. Ask key questions such as:
 - Which attacks are being used by threat actors in current campaigns?
 - Which of those attacks are effective against the business applications deployed in an organization's network?
 - Which of those attacks are capable of bypassing the deployed defenses?
- **Withstand.** How can the business continue to function in an infected environment without failing? How can the organization continue to function without having to suspend critical services to remediate the problem? Assuming that some impact is inevitable, the organization should prepare to operate at 60 percent capacity in order to withstand the breach, which will reduce, but not eliminate, services. It should also segment networks so that a compromised low-priority host cannot infect the rest of the organization and cause system-wide loss.
- **Recover.** Reconstitute and rearchitect the network for a functional restoration or continuation of critical services. Remove the infected portion of the network and assign new resources to replace it.
- **Evolve and Adapt.** Discover why the attack was successful while it is still underway, and then redesign the architecture to withstand similar attacks.

Reading List

Modeling Exploit Evasions in Layered Security. NSS Labs

<https://www.nsslabs.com/reports/modeling-exploit-evasions-layered-security>

The Targeted Persistent Attack (TPA) – The Misunderstood Security Threat Every Enterprise Faces. NSS Labs

<https://www.nsslabs.com/reports/targeted-persistent-attack-tpa-misunderstood-security-threat-every-enterprise-faces>

Top 20 Best Practices to Help Reduce the Threat of the Targeted Persistent Attack. NSS Labs

<https://www.nsslabs.com/reports/top-20-best-practices-help-reduce-threat-targeted-persistent-attack>

Are Asset-Centric Security Models Outdated? NSS Labs

<https://www.nsslabs.com/reports/are-asset-centric-security-models-outdated>

Protect Information, Not Devices. NSS Labs

<https://www.nsslabs.com/reports/protect-information-not-devices>

Future of Computing: Problems Ahead. NSS Labs

<https://www.nsslabs.com/reports/future-computing-problems-ahead-0>

Correlation of Detection Failures. NSS Labs

<https://www.nsslabs.com/reports/correlation-detection-failures>

Incident Response Part 1: Does It Matter, Or Was It Just Noise? NSS Labs

<https://www.nsslabs.com/reports/incident-response-part-1-does-it-matter-or-was-it-just-noise>

Incident Response Part 2: Breach Found. Did it hurt? NSS Labs

<https://www.nsslabs.com/reports/incident-response-part-2-breach-found-did-it-hurt>

Contact Information

NSS Labs, Inc.
206 Wild Basin Rd
Building A, Suite 200
Austin, TX 78746 USA
info@nsslabs.com
www.nsslabs.com

This analyst brief was produced as part of NSS Labs' independent testing information services. Leading products were tested at no cost to the vendor, and NSS Labs received no vendor funding to produce this analyst brief.

© 2014 NSS Labs, Inc. All rights reserved. No part of this publication may be reproduced, copied/scanned, stored on a retrieval system, e-mailed or otherwise disseminated or transmitted without the express written consent of NSS Labs, Inc. ("us" or "we").

Please read the disclaimer in this box because it contains important information that binds you. If you do not agree to these conditions, you should not read the rest of this report but should instead return the report immediately to us. "You" or "your" means the person who accesses this report and any entity on whose behalf he/she has obtained this report.

1. The information in this report is subject to change by us without notice, and we disclaim any obligation to update it.
2. The information in this report is believed by us to be accurate and reliable at the time of publication, but is not guaranteed. All use of and reliance on this report are at your sole risk. We are not liable or responsible for any damages, losses, or expenses of any nature whatsoever arising from any error or omission in this report.
3. NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY US. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, ARE HEREBY DISCLAIMED AND EXCLUDED BY US. IN NO EVENT SHALL WE BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, PUNITIVE, EXEMPLARY, OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
4. This report does not constitute an endorsement, recommendation, or guarantee of any of the products (hardware or software) tested or the hardware and/or software used in testing the products. The testing does not guarantee that there are no errors or defects in the products or that the products will meet your expectations, requirements, needs, or specifications, or that they will operate without interruption.
5. This report does not imply any endorsement, sponsorship, affiliation, or verification by or with any organizations mentioned in this report.
6. All trademarks, service marks, and trade names used in this report are the trademarks, service marks, and trade names of their respective owners.