GLOBAL THREAT INTELLIGENCE REPORT

2014

NTT Group

Table of Contents

THE SHIFTING THREAT LANDSCAPE

The Shifting Threat Landscape

Cybersecurity threats are not static. They are reality, and are actively working against the organization's infrastructure, applications, information and people. To face this change in the threat landscape, organizational security must evolve to include fast, nimble and active responses; not the traditional concepts of the past.

Organizations large and small have grown beyond traditional physical boundaries, reaching across national borders to invoke resources and capabilities globally.

- "Good enough" basic security is enabling the threat actors to tune and automate attacks to maintain constant pressure on the perimeter of the organization until it is compromised.
- "Good enough" needs to be replaced with "well executed" and "basic security" needs to serve as a strong foundation for "advanced capabilities" to meet this constant pressure.
- Fighting back with traditional solutions is a failing strategy as attackers pour resources into circumvention, and skip over the defenses to exploit the more lightly defended interior.

Security Goes Beyond Geographical and Organizational Borders

Organizations large and small have grown beyond their traditional physical boundaries, reaching out of the organization and across national borders to tap resources and capabilities across the globe. The rise of borderless capabilities overwhelms and breaks the implementation of traditional security controls.

- Managing the perimeter is the new paradigm. While the traditional perimeter was between "us" and "them" it has changed to include our partner or team for today which will be different than the one for tomorrow.
- Our trust model needs to be inverted; where trust is tied to the perimeter of the project, data or team and no longer tied to a single organization or individual resource.

Employee Engagement is Vital

Employees around the world demand anywhere, any device access to resources and capabilities. The organization must respond with a different thinking and a security model which is not tethered to the concept of protecting devices but rather to protecting functionality and data.

Employees around the world demand anywhere and on any device access to resources and capabilities.

- Threats move with data and capabilities. Threats are no longer confined to an exterior perimeter. In fact, there is functionally no inside and no outside. Instead, security becomes about targets from any vector; interior, exterior, on premise, in the cloud and in mixed hybrids of every environment.
- Without getting the basics of security right, organizations do not have the ability to enable the advanced, dynamic capabilities needed.
- In an engaged world, static security which is based on hard assets cannot succeed. New security practices have evolved to make dynamic security, tracked to the object, the new standard.

Security Must be Built into Applications

Application development and acquisition is no longer just about putting a front end on a business capability; it has become a bearer and, in some cases, source of corporate risk. Organizations are judged and monetized (or liable) for how well their security works in an agile multi-location public and privately managed business model.

- It's not just how well the application is secured; but how well it is developed, architected, configured and maintained over time which matters.
- The method of development, the infrastructure (public, private, hybrid, other) and effectiveness of security process management (across many hands) have all become keys to maintaining the ongoing risk reduction and application protection.

EXECUTIVE SUMMARY

Executive Summary

The Shifting Threat Landscape

Security is the counterbalance to risk. As the world fluctuates and evolves, so too does the pendulum of risk vs. security. Security is not just about doing the right thing, but also about constant improvement, vigilance and the response to the dynamics of change. Ed Snowden's NSA leak is a prime example of an exploitation of the dynamics of this landscape and abuse of the long-term standard of "good enough" security. Other recent high profile breaches are even better examples of the failures of being "good enough."

Ed Snowden's NSA leak is a prime example of an exploitation of the dynamics of this landscape and abuse of the long-term standard of "good enough" security. Other recent high profile breaches are even better examples of the failures of being "good enough."

Target, for instance, had a reputation for being a well-respected organization, but, as a result of their 2013 breach, has acknowledged they failed to observe some basic security controls, including maintaining appropriate network segmentation, an active patch management process and an event response process. This highlights the need for strong security fundamentals, and not just technology, but process. Target did not face a basic attack. They faced a carefully planned attack which was executed in a very calculated manner. The attack was designed to evade basic controls. Target was using advanced techniques which successfully identified hostile activity in their network. By Target's own admission, this was a failure in fully using available security capabilities and controls; without the operational security foundation in place, Target was unprepared to act when presented with advanced information. In this Global Threat Intelligence Report, the NTT Group security companies (NTT Innovation Institute, Solutionary, NTT Com Security, Dimension Data, and NTT Data) detail how the shifting threat landscape has just begun impacting the long-term implementation of security as usual and how it will continue to shape the face of corporate information security for the future.

While risk is shaped dynamically, the speed and agility of business, tied to the shift towards high mobility in the organization has dramatically changed the meaning and impact of past security investment. Individuals, both personally and professionally, are consuming, shifting and storing more data objects than ever before. And this data is no longer confined to just the corporate network, but rather spread across vast domains inside the organization, in the private and public clouds, in social media, as well as the massive number of continuously connected, offline and near-line endpoints. Protecting data is no longer confined to the hard outer shell of the organization but into a complex fluid environment which goes beyond geographic and organizational borders.

Streams of data are now embedded in the enterprise context, as applications, users and networks shift information to the point of need, possibly storing it locally in case of future access requests. Tied intimately with this flood of data is the security responsibility and management of the rising value of the information which can be exacted from the data flow. In effect, the demand of the end user has made them the gatekeeper of information security. Applications are not simply a capability used by end users, but rather the security envelope which manages the container of these data streams.

The proliferation of data objects has become an effective and separate living entity, an Internet of Things, and will continue to evolve in the future. With this, the web of the security maintenance environment continues to grow and stretch. Security organizations, as shown in this report, have vast challenges under the existing security operational model to maintain wrappers around data objects especially when external and internal environments are addressed differently.

While businesses drag legacy security along, cyberterrorism and criminal enrichment are carrying the attacks forward at a furious rate. Incidents tied to these activities are no longer isolated security events, but rather represent a shift to long-term infiltration. Security's responsibility is to ensure continuous business operation in vastly different environments than legacy capabilities are designed to manage. It is no longer just about protecting critical infrastructure but also the security and safety of the organization, data and people. Security done right, needs to move to the next level of investment so the basic embedded security fabric is the corporate way of doing things, rather than the ugly stepchild to business as usual.

Key Findings

The Perimeter of One

Embedded throughout this report is a strong case that the traditional network perimeter is eroding and single endpoints are becoming the bastion for active threat mitigation including evidence such as:

- **43% of incident response engagements were the result of malware against a particular end point.** Significant factors in these engagements were missing basic controls, such as anti-virus, anti-malware and effective lifecycle management.
- **Research indicates anti-virus fails to detect 54% of new malware collected by honeypots.** Our research shows 54% of malware designed to take over a compromised systems went undetected by the anti-virus solutions used. 71% of new malware designed to make money or steal information from these systems went undetected as well. This finding supports the premise that simple endpoint solutions, while useful at controlling some malware threats, are not capable of fully defending those endpoints against many modern attacks, and must be augmented with network malware detection and purpose-built solutions.
- **An open environment such as education has the lion's share (42%) of malware events.** This was largely due to the open access model typical of universities and the inability to enforce security controls across thousands of student-owned devices. When left to manage security, the end user becomes the first and last ill-prepared defense.

Application Security is now a Business Capability

Traditionally, application security (Internet and internal) has been seen as an adjunct to effective business capability. As the application is now being carried to the forefront of corporate image, and as a result, security, and the ability to maintain that security, are core to maintaining business trust. Yet application security is often still an organizational weakness. This is demonstrated by:

- **Basic application security is often the leading cause of incidents.** Within this report are details on how most of the costs associated with responding to incidents are due to missing or improperly functioning basic controls, inadequate planning and lack of formal training.
- **Cost of a single unsanitized field during automated SQL Injection attack can cost an organization $196,000.** Organizations must realize the true cost of an incident, including direct and public trust costs, and learn how a small investment could reduce losses by almost 95%. Read the "Massive Data Exfiltration via SQL Injection" case study included in this report to see how one unsanitized field on a web form cost an organization over $196,000.
- **77% of organizations supported during incident response activities had no incident response plan.** It is disturbing most organizations have little to no investment in defining and validating (through effective testing) a plan to help navigate critical incidents and minimize damage to their systems, their customers and their brand.
- **Distributed Denial of Service (DDoS) attacks accounted for 31% of incident response engagements.** Most organizations fail to realize the impact a DDoS attack can have, or believe they will not be targeted. They neglect to budget for proactive controls to mitigate DDoS attacks. Scrambling for budget, purchasing solutions and obtaining approval to implement controls while a DDoS attack is occurring has proven to be ineffective.
- **Botnet activity accounted for 34% of events observed in 2013.** The primary targets of healthcare, technology and finance, accounted for 60% of observed botnet activity. This reflects how much these industries rely on the use and flow of information and how dependent they are on maintaining application security for business continuity.

31%

Distributed Denial of Service (DDoS) attacks accounted for 31% of incident response engagements.

The tighter the focus on security in an area, the lower the risk in the area

While this may appear to be obvious, it is harder to quantify in the real world. Information security investment metrics do not necessarily directly equate to success. Yet, organizations which step up into very proactive roles have demonstrated reductions when compared to peers who are less focused

- **Payment Card Industry (PCI)-assessed organizations are better at addressing perimeter vulnerabilities.** Organizations performing quarterly external PCI Authorized Scanning Vendor (ASV) assessments have a smaller vulnerability footprint, as well as a faster remediation time (35%), than organizations which perform assessments but without similar regulatory requirements.
- **Mature vulnerability management reduces threat exposure.** Organizations utilizing a Vulnerability Lifecycle Management (VLM) process have a 20% faster remediation time.
- **Organizations lacking mature VLM programs are four times more susceptible to attacks via exploit kits.** Analysis of total vulnerabilities contained in exploit kits shows an organization without a VLM program generally has as many as four times as many exploitable vulnerabilities in its environment than an organization with a mature VLM program.

Back to basics is not optional, but a required capability of security organizations

Of all the findings in this report, the clearest finding is basic, repeatable, ongoing security measures are core to an organization's success in meeting immediate security challenges as well as addressing the future direction of information security. This is demonstrated by facts such as:

- **50% of vulnerabilities identified in scan data for 2013 were first discovered and assigned Common Vulnerabilities and Threats (CVE®) numbers between 2004 and 2011.** This indicates a massive gap between the detection and remediation phases of VLM, indicating failure of a basic security control.
- **Organizations are at risk to vulnerabilities in the wild.** Data collected in 2013 demonstrates many organizations are not protected against common vulnerabilities which are included in widely distributed hacking exploit kits, allowing these kits to pose a significant threat to their organizational security. We identify the top 10 vulnerabilities identified in customer environments which are also present in exploit kits. Organization should confirm these issues are fully remediated in their environments.
- **Exploit kits ramp up their capabilities to leverage recent vulnerabilities.** Research indicates exploit kit developers are pruning older exploits and favoring newer ones, as 78% of current exploit kits are taking advantage of vulnerabilities less than two years old. While organizations are getting around to removing, updating or patching older systems, attackers are addressing the eventual change in the environment, knowing a history of poor maintenance practices are likely to repeat

Recommendations

The NTT Group security companies believe there are many opportunities for organizations to improve their threat profile and emphasize advanced techniques to detect, investigate and respond. In the technical details below, NTT Group illustrates many organizations are not keeping up with basic controls, and as threats become more sophisticated, organizations must evolve to address today's trends and attacks.

In the 2014 GTIR, NTT emphasizes advanced techniques to detect, investigate and respond. This report contains information organizations can use to improve their operational security and how proper implementation, management and adjustment of proven controls can help organizations reduce risk by avoiding threats and compress the mitigation timeline, significantly reducing loss exposure.

- **Address the eroding perimeter.** As detailed previously, the threat is shifting and the perimeter is becoming vastly different than traditionally envisioned. To meet the shrinking perimeter, key components which can be incorporated now includes using patch management and anti-virus programs to maintain a last defense, however small the perimeter becomes. When combined with accurate asset inventories and escalation SLAs, these basic controls can help organizations limit risk from well-known vulnerabilities and the most common attack vectors.
- **Use effective patch management to protect against real-world threats.** Effective patch management is not easy, and timely installation of every patch on every system is often impractical. Additionally, attackers often have a long lead time on vendor remediation. Organizations need to be aware of issues being exploited in the real world, and need to ensure they are prioritizing countermeasures against these exploits.
- **Define and test incident response.** Too many organizations have untested, immature or non-existent incident response programs. This makes them unprepared for the inevitable attack, especially in an evolving world where organizations are no longer safe behind strong and secure exterior perimeters. Appropriate incident response is critical to minimize the impact of security breaches. All organizations need to document, test and maintain effective incident response procedures.
- **Take advantage of new technologies and techniques.** While patch management and anti-virus are critical components of a security program, they will continue to erode in coming years. The speed of exploit weaponization is increasing and may surpass an organization's ability to respond quickly and effectively (if it has not already). New technologies include capabilities such as application isolation techniques, micro VMs, sandboxing and machine learning. These technologies focus on application control and isolation, incident containment and rapid detection via behavioral analytics, are likely to grow in importance. These technologies assume the perimeter will fall and compromise is inevitable, and while some preventive techniques can help, the best defensive approach is to limit exposure and detect (and respond to) incidents quickly.

Global Data Analysis and Findings

This section presents an analysis of global attack data gathered from NTT Group security companies in 2013. The analysis is based on log, event, attack, incident and vulnerability data from clients and NTT research sources, including honeypots and sandboxes. It summarizes data from trillions of logs and over three billion attacks.

NTT Group gathers security log, alert, event and attack information, enriches it to provide context, then analyzes the contextualized data. This process enables real-time global threat intelligence and alerting. The size and diversity of our client base makes this data representative of the threats encountered by most organizations.

The data is derived from correlated log events identifying attacks based on types or quantities of events. The use of validated attack events, as opposed to the raw volume of log data or network traffic, more accurately represents actual attack counts. This methodology lends credibility to the resulting data. Without proper categorization of attack events, the disproportionately large volume of network reconnaissance traffic, false positives, authorized security scanning and large floods of DDoS actively monitored by Security Operations Centers (SOCs), would obscure the actual incidence of attacks.

Ultimately, the inclusion of data from the various NTT Group security companies allows for a more accurate representation of the threat landscape around the world.

2013 ATTACK ANALYSIS

2013 Attack Analysis

Attacks by Country

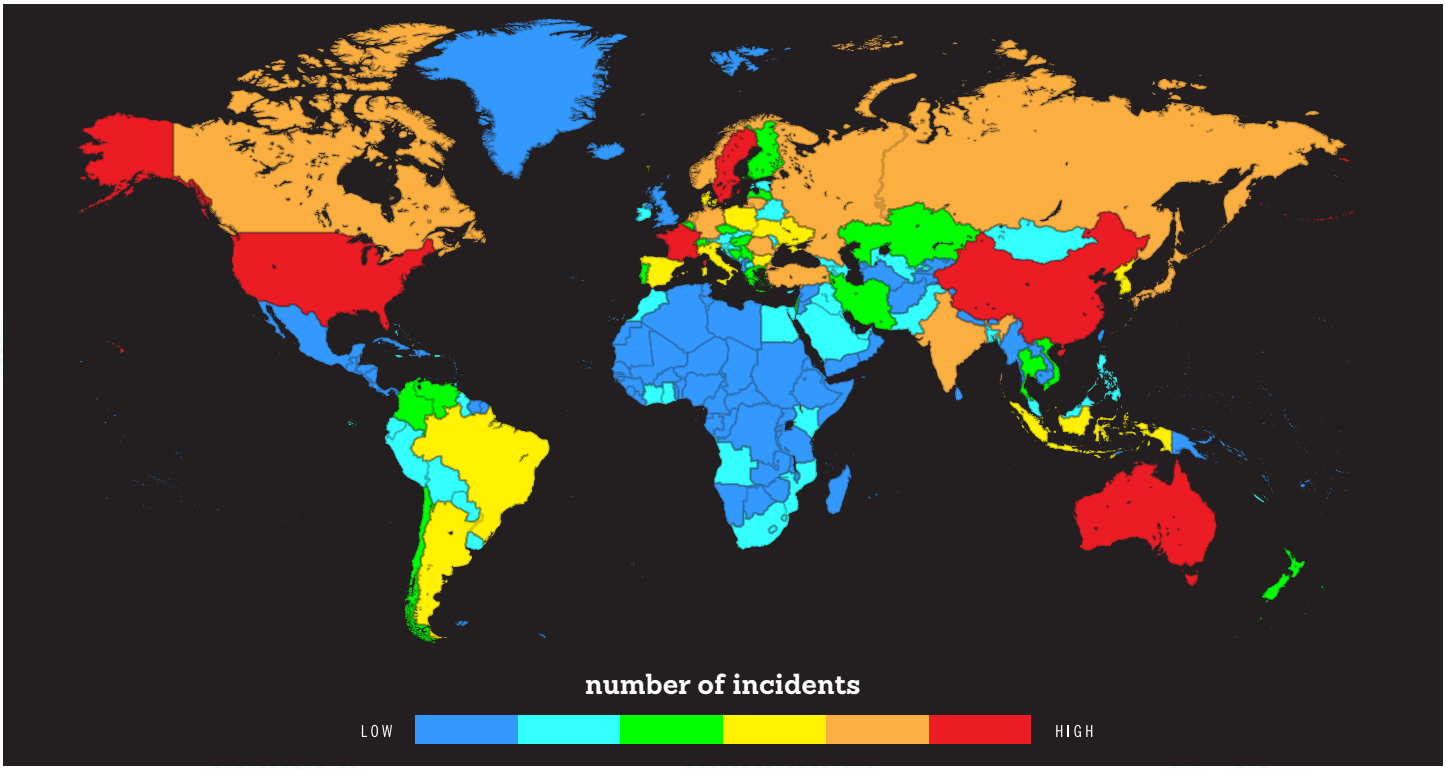

Figure 1: Worldwide Map of Attacks by Country

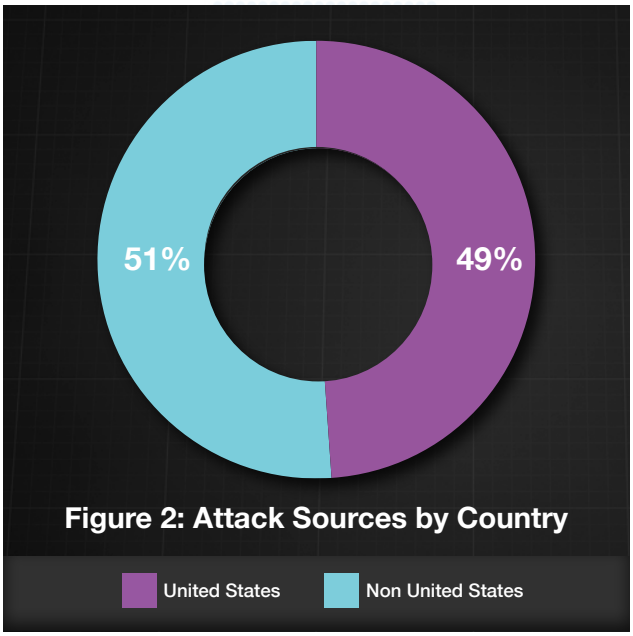

Figure 2: Attack Sources by Country

As shown in Figure 2, 49% of attacks against the NTT client base during 2013 originated from U.S. source addresses. In general, attacks within regions appear to be growing, making geographic blocking ineffective.

Deeper analysis indicates attacks appearing to come from U.S. addresses often originate in other countries, but use U.S. IP addresses as a way to hide their actual location. We have observed this target localization strategy in other parts of the world as well, where attackers establish a point of presence in the same country as their target, and then use this as a launching point for attacks and exfiltration of data. As documented in the NTT Group Q4 2013 SERT Threat Intelligence Report[1], for U.S. based attacks, we continue to observe frequent use of highly-regarded hosting providers such as Amazon Web Services® and GoDaddy® as temporary data exfiltration drop points.

[1] *http://www.solutionary.com/research/threat-reports/quarterly-threat-reports/sert-threat-intelligence-q4-2013*

Attacks by Industry

Figure 3 represents the distribution of attacks from all sources against the targeted industry vertical. The Finance and Technology industries continue to lead other industries as targets. Figure 4 is included as a reference for the client makeup that contributed to the data used in this report.

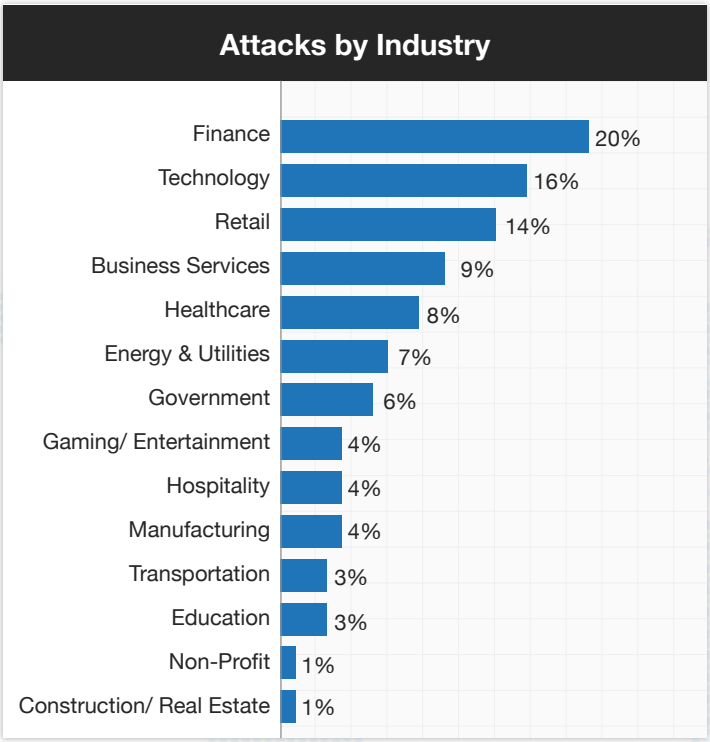

Figure 3: Attacks by Industry

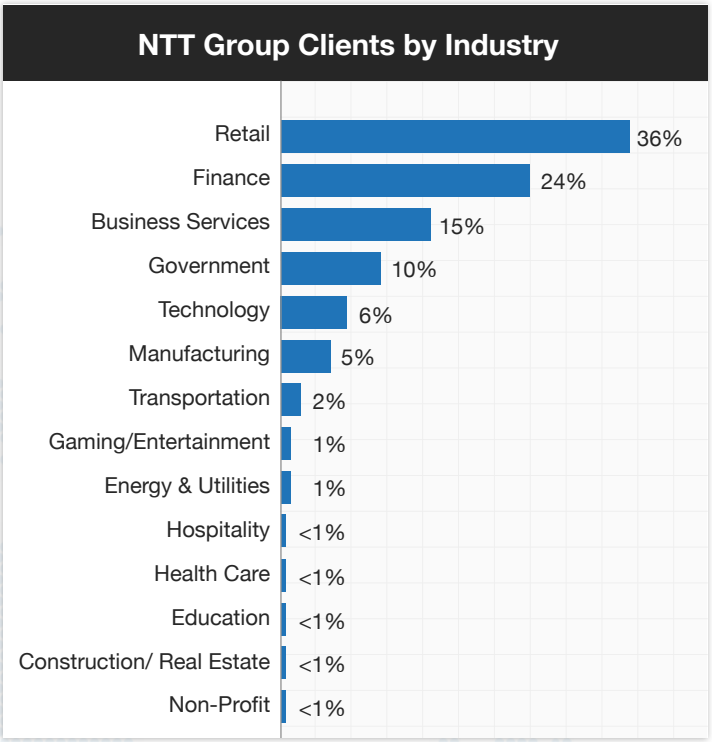

Figure 4: NTT Group Clients by Industry

Attacks by Type

Figure 5 illustrates the distribution of attack types from all verticals and countries. Client botnet activity was the largest type of attack, making up 34% of attacks.

In a botnet attack, the attacker distributes malicious software (malware) to compromise systems. This malware helps the attacker remotely control the distributed systems in directed, automated actions, without the user's knowledge. The attacker's network of controlled systems is called a botnet.

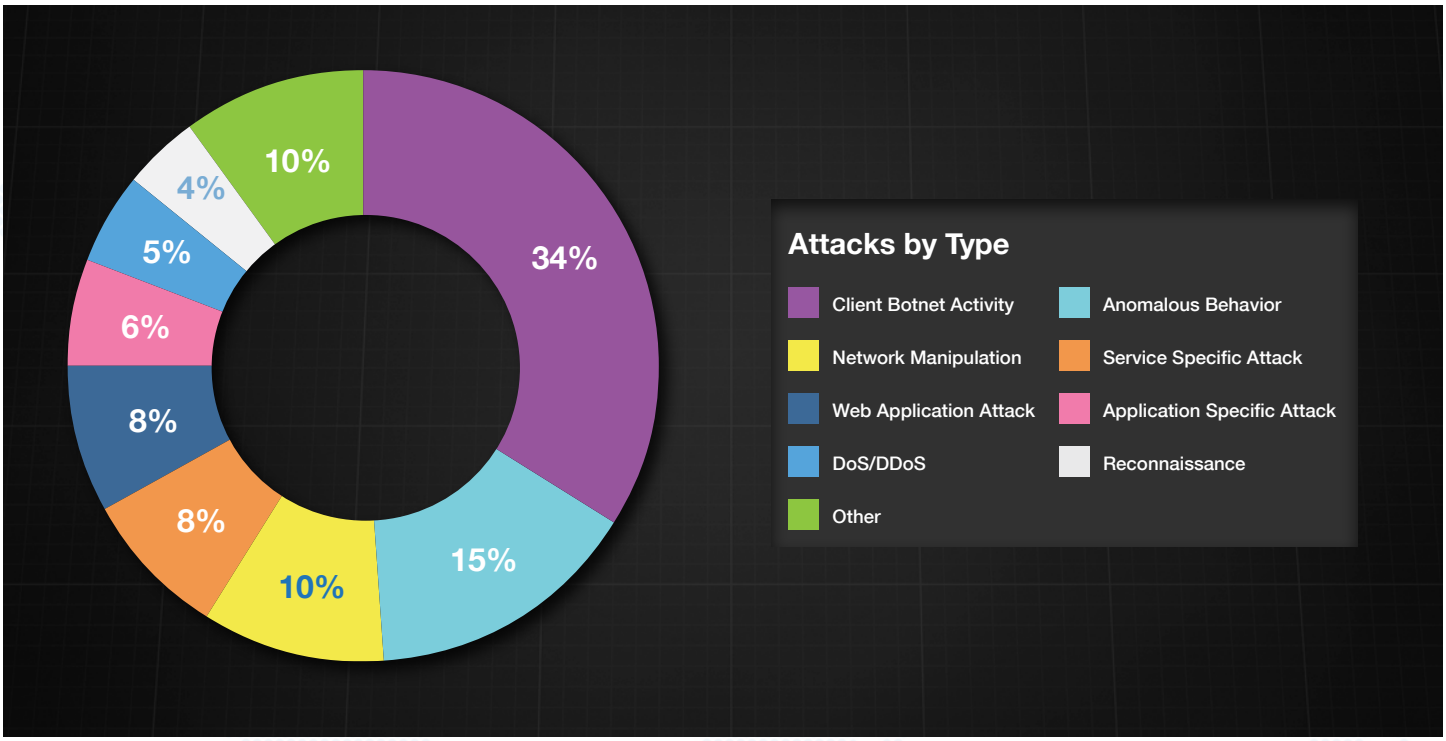

Figure 5: Attacks by Type

This finding can be attributed to the widespread use of bots and Command and Control (C&C) activities, along with the use of botnets to support DDoS attacks. Anomalous behavior, such as high volumes of dropped packets and traffic not following RFC guidelines, is the next highest category, representing 15% of total events. Network manipulation tactics, such as DNS attacks and IP address spoofing was third highest with 10%. Other than the types displayed in the chart, no other single attack type accounted for more than 2% of the identified attacks.

While malware represents only 2% of the attacks identified in 2013, modern malware is most often used as a component of more advanced attacks and weaponized tool sets. This results not only in infection of the target host, but also in other types of breaches when the malware opens the door to exploit kits, lateral attacks within a network and further spread of the malware via spam forwarders and compromised web sites.

Client Vulnerabilities

Vulnerability data for 2013 is composed of a wide range of scanning data from multiple scanning vendor products (Qualys, Nessus, Saint, McAfee, Rapid7, nCircle and Retina) and across multiple organization verticals and sizes.

Top 10 External Vulnerabilities	Percentage of All External Vulnerabilities	Top 10 Internal Vulnerabilities	Percentage of All Internal Vulnerabilities
Outdated Apache Tomcat Server	18%	Missing Security Updates for MS Windows	6%
Outdated Apache Web Server	13%	Adobe Reader Vulnerabilities	6%
Cross Site Scripting Vulnerability	10%	Outdated System Management Consoles	6%
Outdated PHP Version	9%	Oracle Java SE Critical Patch Update	5%
Server Side Includes - Injection Vuln	5%	Outdated Apache Web Server	3%
Web Clear Text Username/ Password	3%	Outdated Adobe Reader and Acrobat	2%
Vulnerable 3rd Party Apache Plugins	3%	Outdated OpenSSH Version	2%
Outdated OpenSSL	2%	Sun Java J2SE 1.4.2 < Update 18	2%
Cookie without HTTPOnly attribute set	2%	Outdated Java Runtime Environment	2%
Cross-site Request Forgery Detected	2%	Internet Explorer Remote Code Execution	2%

Figure 6: Top 10 Vulnerabilities in 2013

Companies tend to exhibit M&M® candy-style architectures, where externally-facing systems show a more secure profile (hard shell exterior), while internal systems are less secure (soft gooey interior). This is not unexpected, given typical budgets and priorities for security controls. The top 10 external and internal vulnerabilities identified in 2013 are presented in Figure 6.

Figure 7 shows the distribution of vulnerabilities by type, as identified in 2013. This data shows most organizations are still struggling with security basics. The majority of vulnerabilities identified are due to patch management, firewall and application configuration, and other basic security practices.

Scan results for the top 10 vulnerabilities echo the themes discussed earlier. Organizations must improve their patch management process as well as application configuration

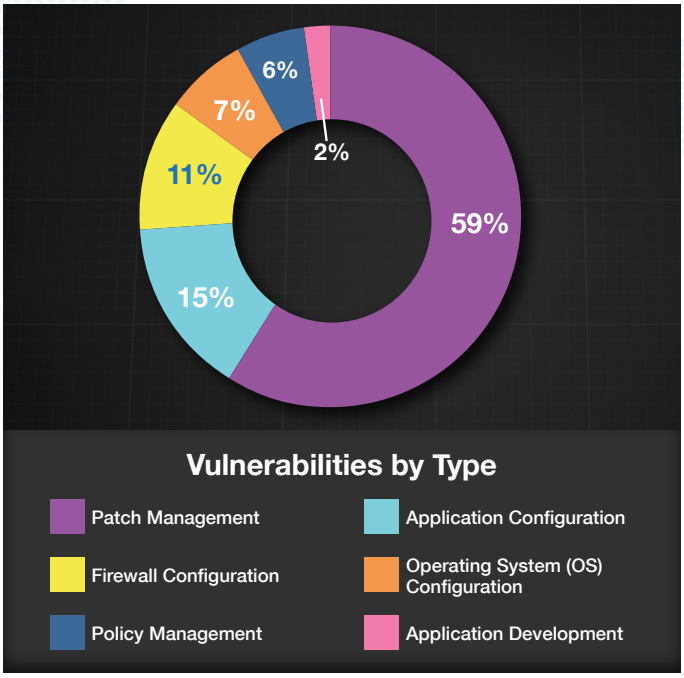

Figure 7: Vulnerabilities by Type

and development issues as tracked by the Open Web Application Security Project (OWASP), including cross-site scripting, SQL Injection and cross-site request forgery.

Exploit Kit Trends

Looking at the CVEs exploit kits are leveraging, it is clear to see how they align with the typical M&M security model. Most malware is intended to be delivered directly to end-user devices, bypassing external network defenses. If malware can be delivered internally, the target organization is less likely to detect and mitigate.

CVE	Percentage of Exploit Kits Found In[2]	Affected Technology
CVE-2012-1723	37%	Oracle Java SE
CVE-2013-2423	35%	Oracle Java SE
CVE-2013-1493	26%	Oracle Java SE
CVE-2013-2471	21%	Oracle Java SE
CVE-2013-2460	19%	Oracle Java SE
CVE-2013-2551	16%	Microsoft Internet Explorer
CVE-2013-2465	16%	Oracle Java SE
CVE-2013-0431	14%	Oracle Java SE
CVE-2013-2463	14%	Oracle Java SE
CVE-2013-3896	7%	Microsoft Silverlight

Figure 8: 10 Most Common Client Vulnerabilities in Exploit Kits

The top ten matches between the most common vulnerabilities and the CVEs in exploit kits are presented in Figure 8.

Organizations should consider this an area on which to focus active patch management and configuration efforts. Since these vulnerabilities are common and are being widely exploited, organizations need to ensure they are actively addressing these vulnerabilities.

Exploit kits also appear to be better at utilizing more recent vulnerabilities than we have previously observed. Figures 9 and 10 show the distribution of CVEs by year included in exploit kits, from 2012 and 2013. In 2012, we observed the average age of vulnerabilities included in exploit kits was slightly less than two years old, while in 2013 the average vulnerability age was just over a year old. In both years, kits included exploits from as old as 2006. This indicates that attackers are growing their sophistication and ability to rapidly update exploit kits before organizations have the chance to react.

[2] *This chart includes data from http://contagiodump.blogspot.com/2010/06/overview-of-exploit-packs-update.html. This site is a valuable malware and exploit kit resource.*

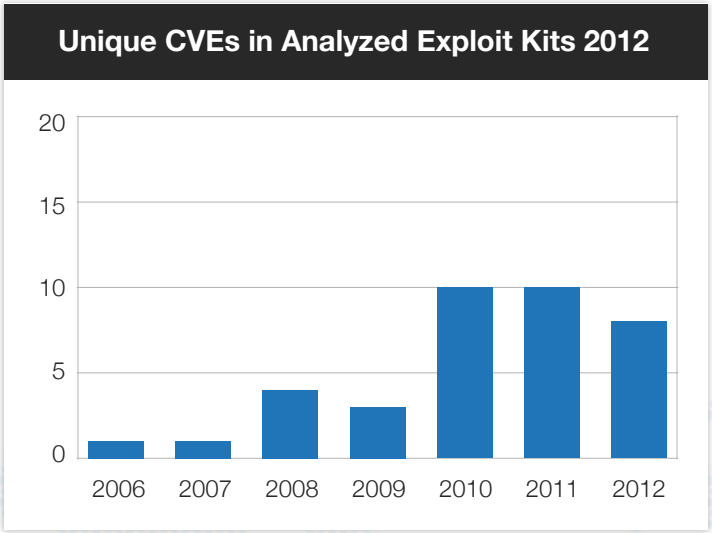

Figure 9: Unique CVEs by Year in 2012 Exploit Kits

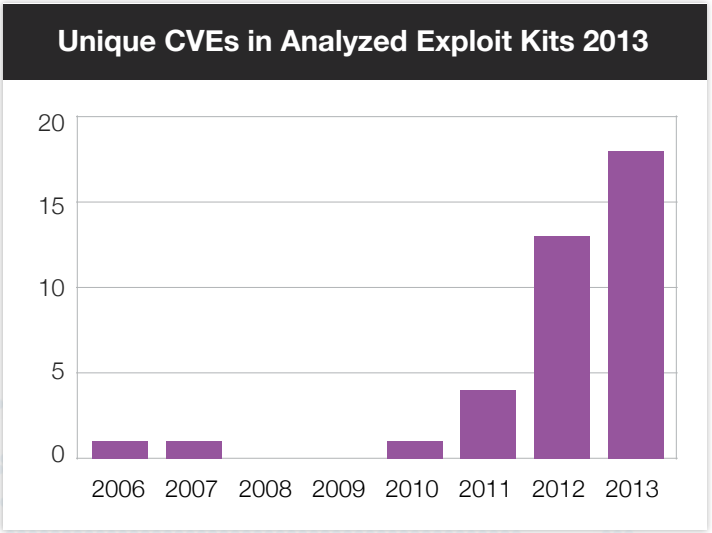

Figure 10: Unique CVEs by Year in 2013 Exploit Kits

Today's complex threat landscape highlights the need for organizations to increase awareness of active patch management, and to prioritize patching of issues posing the greatest risk.

KEY THREAT OVERVIEW AND CASE STUDIES

Key Threat Overviews and Case Studies

During 2013, several key threats were recurring themes. Malware, botnets and web application attacks can have a profound effect on organizations. Malware is becoming more complex, and has become a key component of other attacks. Attackers are using malware to spread influence and control throughout a victim's network, and to support extended threats such as botnets. Botnets have spread rapidly, taking over internal and external systems and consuming organizational resources. Attackers continue to focus on web application attacks because these attacks provide a gateway to critical resources such as client databases.

The most effective way to talk about the impact these issues have on an organization is to use real-life examples. Our case studies are derived from NTT Group investigations in 2013, with actual observations and recommendations. These studies include identification of security gaps which led to the investigation, the controls which enabled the organization to react effectively to the incident, or the lack of controls which made reaction inefficient. These studies also include recommendations that would have allowed the organizations to react more efficiently and compress mitigation timelines.

The descriptions of these key threats, and the case studies that support them, clarify the impact (good and bad, depending on maturity) that detective, investigative and response capabilities can have on an ongoing incident. We encourage you to consider the ways in which these case studies reflect the security posture of your organization, and to consider if these recommendations may be appropriate for your situation.

Malware

In 2013, NTT Group's research focused heavily on malware analysis. Beyond the initial infection, malware facilitates other attacks across networks, and extends the reach of botnets such as ZeroAccess. NTT Group gathered samples of malware from a wide range of sources, including analytics and correlation data from NTT Group security companies' security platforms, incident response investigations, malware repositories, malware feeds, interactions with clients and privately maintained honeypot networks.

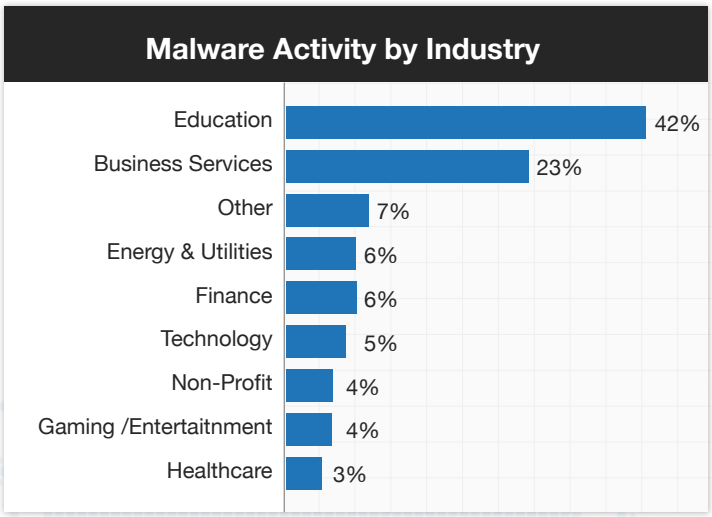

Figure 11: Malware Activity by Industry

Malware by Industry

The 42% of malware events contained within the education segment makes sense from an organizational standpoint. Universities and other educational institutions have a large body of users connected to a public network with personal systems and a culture which promotes making information readily available. In this challenging environment, some organizations manage well, but the requirement to enforce security controls in a relatively unenforceable environment represents a significant risk to information security. These environments include student-owned computers and mobile devices with a wide variety of operating systems, plus different degrees of endpoint protection and end-user awareness, as well as patch levels. This represents a significantly more complex environment than the typical corporate technology ecosystem.

Anti-Virus Detection Rates

NTT Laboratories evaluated all acquired malware samples by systems running 11 different commercial or free anti-virus engines. As shown in Figure 12, the common anti-virus software used failed to detect as much as 54% of malware samples acquired via honeypot systems. This includes malware gathered by crawling for vulnerable and compromised websites with our proprietary honeypots which were pretending to be client PCs. The malware acquired by the honeypots is then run in our proprietary sandbox, pretending to be an infected host. The malware may be instructed to download additional applications like key loggers, or software to steal Internet banking and credit card information. The malware downloaded after the initial infection is generally more sophisticated and difficult to detect. This is proven by the fact the anti-virus detection rate for this type of malware drops even more, with as much as 71% of such malware avoiding detection.

Mass-distributed malware hides its presence from anti-virus software using advanced techniques ranging from embedded encryption to dynamic variant generators, which make subtle changes to prevent signature-based detection. Additionally, malware authors develop new variants on a frequent basis and test their code against the same anti-virus software organizations use, fine-tuning the malware until the detection mechanisms fail. These factors prevent anti-virus software from being completely effective in preventing infections by themselves. However, organizations should not ignore a solution as straightforward as anti-virus since it can still reduce these threats by almost 50%.
The following case study describes the impact of a malware infection on a client who was not ready to manage an outbreak.

ANTI-VIRUS DETECTION

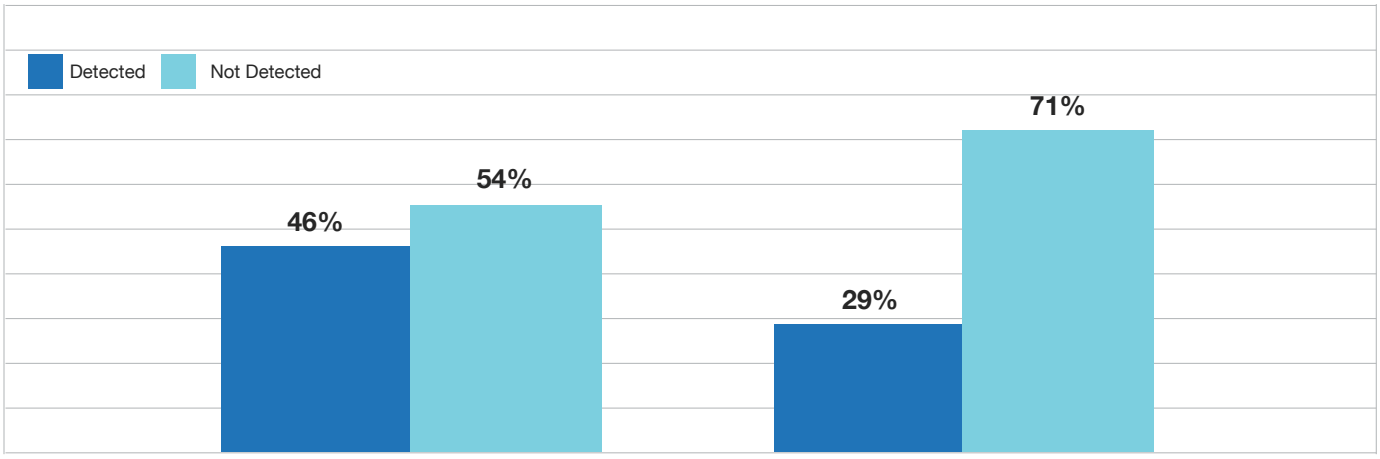

Figure 12: Anti-Virus Detection Rates

…organizations should not ignore a solution as straightforward as anti-virus since it can reduce virus threats by almost 50%.

Case Study:
Administrator Releases a Worm

CLIENT DOSSIER

Industry Vertical:	Government
Location:	Australia
Employee Count:	< 100
Financial Transactions (USD/Year):	$5 – 15 Million
Ownership:	Government

ATTACKER DOSSIER

Location:	Russia
Primary Targets	Targets of Opportunity
Prefered Attack Methods	Malware Distribution

CAPABILITY RATING

DETECTIVE

INVESTIGATIVE/RESPONSE

Case Study Moral:

Even a good security program can be disrupted by poor incident response.

Overview:

In July of 2013, NTT Group discovered Dorkbot malware in the XYZ Government Group[3] (XGG). Rapid detection and isolation was hampered by confusion and incomplete remediation capabilities.

[3] *Organization names have been changed in case studies.*

Timeline of Events

Date	Event Activity	Indicators
Day 0	Initial server compromise	**Undetected:** Active, updated anti-virus would likely have detected and prevented infection on the initial device **Detected:** Alerts triggered on anomalous traffic initiated manual analysis
Day 1	Command and control activities initiated by original infection	**Detected:** Rules for anomalous egress traffic reported alerts
Day 1	Malware dropper installed with application hooking functions	**Undetected:** Active, updated anti-virus would likely have detected and prevented infection **Undetected:** The additional infection was not observed by XGG
Day 1	Failed attempts at administrative login results in account lockouts	**Detected:** Logging reported failed authentication attempts
Day 1	Malware begins probing for open ports to the internet	**Detected:** SOC analysts identified anomalous network traffic
Day 4	Client *claims* availability "restored"	**Undetected:** Server Zero remains infected, along with other unidentified system infections
Day 87	Final restoration of systems and complete removal of infection	**Detected:** Normal baseline traffic observed after complete analysis and restoration of known good backups

Figure 13: Timeline of Events - Dorkbot Worm

Description of Event

In July of 2013, NTT Group detected anomalous activity within XGG. Initial investigation revealed alerts for various types of anomalous traffic, including:

- High volume of blocked packets to a limited number of destinations
- Connections using abnormal ports
- Connections to a known C&C server and malicious domains
- Worm-like propagation over ports not used for normal communications

A "Command and Control" (C&C) server is a centralized server used to control botnets and other malware. Remote malware reports status to the C&C server, and in turn, the C&C server sends additional commands to the remote malware. Remote malware can also upload captured data to the C&C server, or to another server as directed.

The Security Operations Center (SOC) analyst determined that in the context of XGG's business, the destination domains were unusual, and verified egress traffic was not associated with any authorized application in the environment. Forensic analysts determined an authorized administrator had unintentionally loaded and executed an infected file from a USB device onto a server not being directly monitored by NTT Group. The worm connected to a C&C server, uploaded user credentials

and downloaded a malware dropper, which continued to add and conceal other malicious packages. Then the worm initiated the C&C connection over IRC, using parameters obfuscated with base64 encoding. The malware traffic was destined for China and then attempted connections to Ukraine. SOC analysts decoded the communication to discover internal user credentials, computer name, system type, date, locale and a unique bot identifier. Attackers had successfully injected system processes on the infected server (including explorer.exe and vmtoolsd.exe) with malicious code. The malware was successful because XGG had not completely deployed anti-virus solutions on internal servers prior to this attack. As mentioned previously, as many as 54% of viruses are not detected by anti-virus software, but analysts determined this malware was a member of the Dorkbot family which would likely have been detected with current signatures. Forensic analysis also confirmed the worm had replicated itself onto other networked servers.

A "malware dropper" is a program designed to evade anti-virus solutions, download additional malware packages and install applications without user interaction.

XGG initiated incident response quickly, but had not yet implemented all planned security controls, and had not fully implemented or tested their incident response plan. XGG staff's lack of awareness of response practices contributed to confusion of responsibilities and a lack of ownership of mitigation efforts. Poor communication led directly to an incomplete and inefficient recovery process, delaying initial containment of the worm by several days. Management insisted the main priority in the response was continued system availability, instead of successful removal of the malware. During this extended timeframe, additional systems were infected and XGG experienced a serious, prolonged degradation of service.

XGG installed previously missing anti-virus and anti-malware engines on infected systems. This software was subverted by the evasive capabilities of the original infection, which then altered the scans to conceal its presence. This installation process was repeated by different responding groups, using different anti-virus vendors, which removed some malware artifacts without removing the persistent, underlying malware.

Dorkbot is a family of worms that steals user credentials and can use victim devices in coordinated DDoS attacks.

After weeks of struggling with incomplete remediation, XGG management began reimaging the affected systems in order to fully restore services, but the responders were unwilling to take responsibility for the existing situation or for steps necessary to mitigate the attack. XGG found itself having internal arguments about processes and responsibilities while facing an active incident. If the organization had reacted in a timely manner, the attack could have been quickly mitigated after initial identification.

Root Cause

The initial infection was isolated to an administrator who unintentionally uploaded infected software from a USB device. This action was made worse by the fact that many XGG systems had either no anti-virus software or signatures that were out of date. The most significant issue facing XGG was the complete failure of their incident response process, which allowed infections to continue for three months after initial identification. XGG's *inefficient and ineffective response resulted in an incident that cost them $109,000 in measureable response activities, and months of continued malware propagation, degraded service, ongoing troubleshooting and associated mitigation activities.*

Cost of Incident

XGG provided the actual cost of this event. NTT Group can project the expected cost of this event if XGG had implemented basic security controls.

Incident Management Profile	Cost of Incident
The actual cost of investigation, legal support, public relations support, remediation and professional incident support as described.	**$109,000**
If XGG had used a predefined incident response protocol they could have gracefully restored operations in a matter of hours.	**$2,800**
If XGG had fully installed and maintained the anti-virus software that they had already purchased, this incident might have been completely avoided.	**$0**

Figure 14: Cost of Event - Dorkbot Worm

Case Study Summary – Administrator Releases a Worm

Worms and other malware pose a significant threat to any environment, especially when executed by someone with administrator level access.

While network detection methods were in place and discovered the infection quickly, this case presented many opportunities for improved prevention and a more effective mitigation process. The likelihood and impact of infections can be reduced by ensuring administrator access is used only when necessary, by including anti-virus and anti-malware solutions and by having a well-defined and tested incident response process.

Threat Mitigation – Administrator Releases a Worm

In this section we provide guidance on basic and advanced security controls to help protect against malware and related threats, and to enhance incident response capabilities.

XGG's inefficient and ineffective response resulted in an incident that cost them $109,000 in measurable response activities.

Basic Controls – Administrator Releases a Worm

- **Keep anti-virus software up to date.** The Dorkbot malware spread quickly in the XGG environment, partially because not all anti-virus software was up to date, and not all systems were being automatically scanned. Although anti-virus software is not a comprehensive security solution, it can be a valuable part of a layered security approach.
- **Train and educate users.** User training can prevent many incidents, especially malware-related events. In this case, the infection originated with an administrative user who unintentionally installed malware on an XGG server. Organizations should conduct frequent security education sessions for users and administrators to help them recognize malware, phishing and other attacks, as well as their responsibilities in the incident response process.
- **Define and test an incident response process.** After the incident had been reported, XGG staff demonstrated a lack of accountability and responsibility for the incident. As a result, XGG's reaction was haphazard and inefficient. If an organization defines and tests their incident response procedures, they can help facilitate a proper response. Periodic testing greatly increases the chances that an incident response plan will work correctly when needed.

- **Implement proxy servers to control data flow.** Proxy servers add the capability to scrutinize all traffic as it traverses the network. Protocol-aware proxies provide value due to their ability to inspect network communications and compare protocols to those defined by RFCs.
- **Implement egress filtering.** Monitor network communications for aberrant activity across network segments and improper attempts to communicate outside of the organizational network. Egress filtering can detect C&C communications of the botnet.

Advanced Controls – Administrator Releases a Worm

- **Consider a Data Loss Prevention (DLP) solution.** On assets hosting sensitive information DLP technology can be used to track access and potential misuse.
- **Deploy purpose-built network malware detection capabilities.** Malware detection tools can detect complex communication and behavioral characteristics which may otherwise be difficult to track. When paired with an active anti-virus tool, such products can dramatically improve the effectiveness of malware detection.
- **Maintain a whitelist of approved applications.** The security analyst verified XGG did not have any applications authorized to perform the types of communication detected, or to contact the identified destinations. By maintaining an approved software list, an organization can more readily identify unauthorized (or unpatched) software which could pose a risk.
- **Take targeted actions against Dorkbot and associated malware.** An organization can take the following actions to directly address risks associated with Dorkbot and related malware:
    - **Disable autorun and USB capabilities.** While some infections of this family spread via instant messaging programs and social media, one distribution method involves infected USB flash drives. An organization can prevent infected autorun files from executing by disabling the autorun feature or by disabling USB ports for storage devices.
    - **Disable or monitor for unauthorized IRC traffic.** Many infections use the IRC protocol over ports 6660-6669 to communicate with a C&C server. This is likely to change, but at the time this report was written Dorkbot communicates over IRC with the following domains:
        - av.shannen.cc
        - lovealiy.com
        - shuwhyyu.com
        - syegyege.com
    - **Whitelist authorized binaries.** An organization can explicitly whitelist authorized binaries across a Microsoft Windows® network, using the AppLocker feature, with an implicit “Deny All” rule for anything not explicitly allowed. This configuration will prevent execution of new strains of malware, potentially unwanted applications (PUA) and other unauthorized programs.
    - **Monitor CPU/GPU utilization on supported hosts.** Dorkbot and other malware variants are increasingly installing cryptocurrency (e.g., Bitcoin, Litecoin) *miners*. Mining software uses all spare CPU and GPU cycles to process cryptographic hashes, in return for shares of these digital currencies. The result is consistently high utilization of system resources, increased electricity consumption for your organization and profits for the criminals.

Botnets

Botnet C&C Activity by Source Country

Client botnet activity was the largest attack category identified in 2013. Figure 15 shows the United States accounted for 42% of all identified botnet C&C traffic. No other country accounted for more than 4% of the identified botnet traffic.

As shown in Figure 16, botnet attacks by industry differed slightly from total attacks by industry. Healthcare accounted for 21% of botnet attacks.

ZeroAccess Supernodes

Case Study: ZeroAccess Botnet describes a botnet infection and summarizes research that followed. Further research on ZeroAccess supernodes, the infected devices that form the peer-to-peer backbone of the botnet, revealed a total of 360,000 unique IP addresses acting as supernodes in the ZeroAccess network. Like most other threats we see, Figure 17 shows the U.S. included the highest percentage (23%) of the ZeroAccess supernodes.

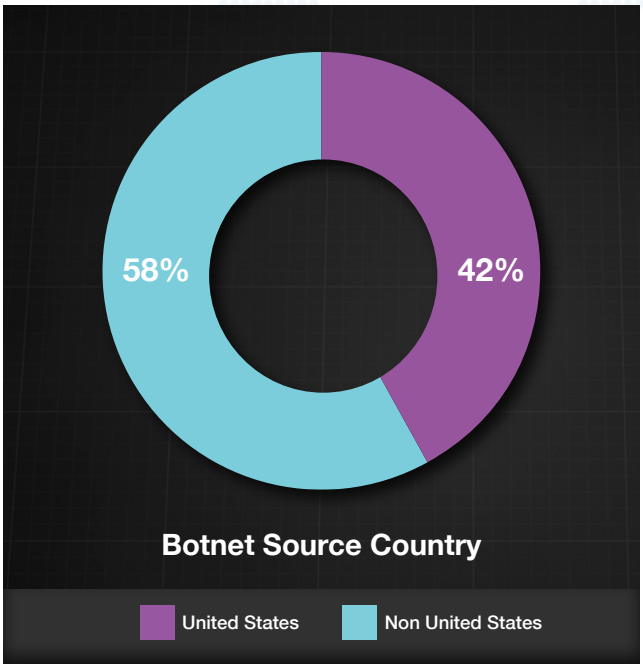

Figure 15: Botnet Source Country

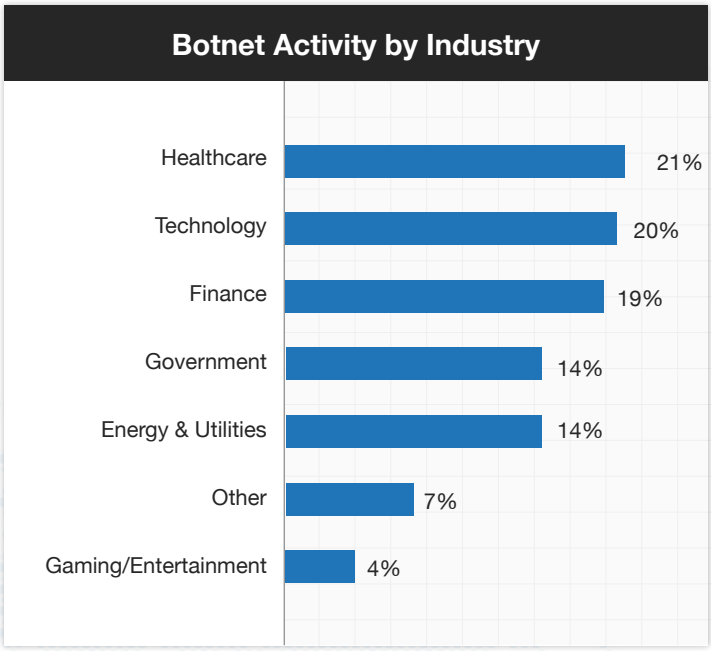

Figure 16: Botnet Activity by Industry

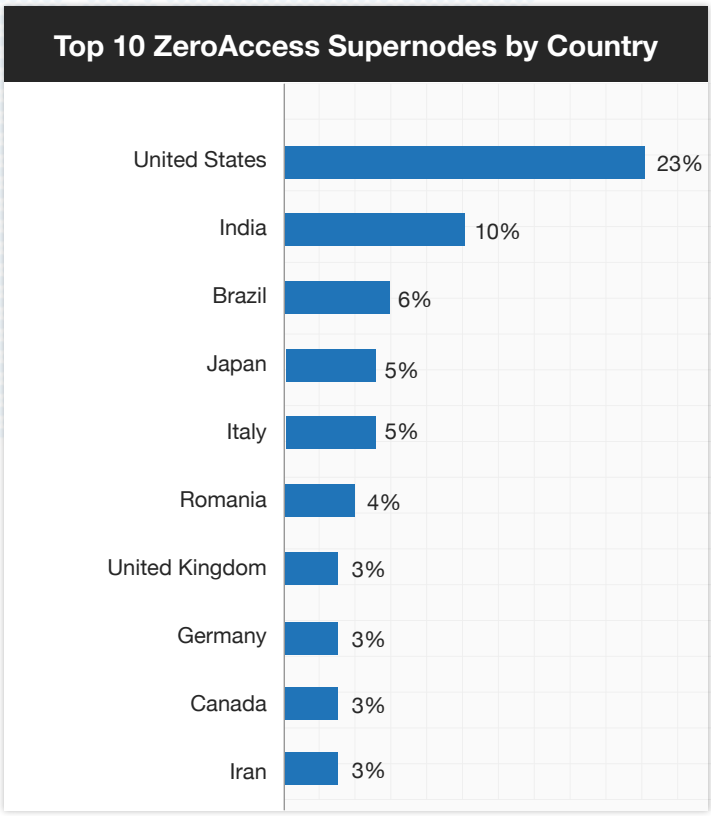

Figure 17: Top 10 ZeroAccess Supernodes by Country

Proper preventative measures, such as active patch management and real-time virus scanning, are still effective methods to protect against this type of widespread threat. Advanced purpose-built, anti-malware solutions can also be very effective at detecting and protecting against botnets.

The following case study describes an organization that detected and responded to the ZeroAccess botnet even though basic security measures failed to stop it.

Case Study:
ZeroAccess Botnet

CLIENT DOSSIER

Industry Vertical:	Information Technology
Location:	Northern Europe
Employee Count:	< 10,000
Financial Transactions (USD/Year):	$2 Billion
Ownership:	Privately Held

ATTACKER DOSSIER

Location:	Germany
Primary Targets	Windows System
Prefered Attack Methods	Botnet Management

CAPABILITY RATING

DETECTIVE

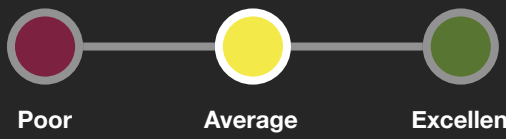

INVESTIGATIVE/RESPONSE

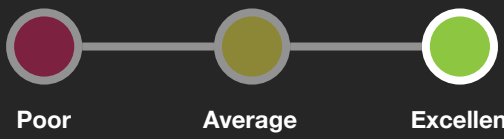

Case Study Moral:

Rapid response can minimize the impact of the lack of basic controls.

Overview:

In June 2013, an analyst for NTT Group discovered indications of the ZeroAccess botnet in XYZ's infrastructure. Efficient incident response within XYZ rapidly isolated the infection.

Timeline of Events

Date	Event Activity	Indicators
Day 0	Initial server compromise	**Undetected:** Active, updated anti-virus would likely have detected and prevented infection **Detected:** IDS alerts triggered manual analysis
Day 0	ZeroAccess supernode identified	**Detected:** Analyst observed inter-node communications
Day 1–5	ZeroAccess signature development and rollout	**Detected:** ZeroAccess communications, command, response and IP addresses
Day 1–8	Anti-virus scanning and system restoration	**Detected:** Systems showed quarantined malware or were rebuilt as "clean"
Day 5–20	ZeroAccess clone monitoring	**Detected:** Over 360,000 worldwide supernodes

Figure 18: Timeline of Events - ZeroAccess Botnet

ZeroAccess Botnet Architecture

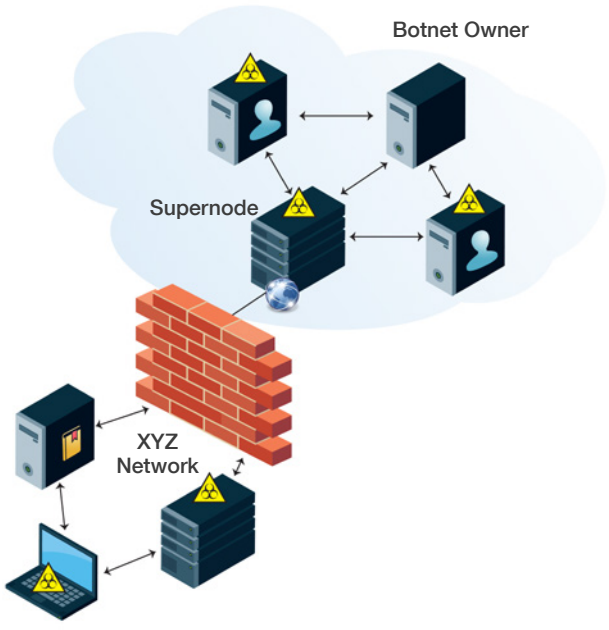

Figure 19: ZeroAccess Example Architecture

Description of Event

In June 2013, NTT Group investigated an IDS event indicating anomalous DNS activity within the XYZ environment. The analyst realized the traffic, while occurring over port 53, was not DNS traffic but resembled Spotify, Skype and other peer-to-peer (P2P) file sharing applications which were not used by this client. The analyst recognized the traffic as ZeroAccess and identified the host as a ZeroAccess *supernode*. The server was communicating with other infected hosts in the XYZ network, using DNS ports as well as UDP port 16464. ZeroAccess is a kernel-mode rootkit that assumes full control of the machine by adding it to the ZeroAccess botnet, then monetizes the asset by downloading additional malware. Supernodes form the backbone of this botnet, providing other infected hosts with information about the ZeroAccess network, and communicating inside and outside an infected network. Figure 19 shows the basic structure of the ZeroAccess network architecture. (Infected nodes not exposed to the Internet simply remain ZeroAccess nodes and support general botnet activities.) NTT Group partnered with XYZ to ensure infected systems were isolated and eliminated in a timely manner. XYZ completed their own remediation using anti-virus, restoring from backups, or rebuilding systems.

A security analyst observed IDS vendor signatures had not detected ZeroAccess traffic, so they performed detailed malware analysis on a sample. This analysis included:

A supernode forms part of the backbone of the ZeroAccess botnet and serves other infected hosts with information.

- Decoding commands used in communication between malicious nodes in the ZeroAccess peer to peer network.
- Examining data flow in response to received commands.
- Mapping IP addresses of nodes with which client systems were communicating.

Based on this analysis, NTT Group developed and deployed custom IDS signatures to detect the initial phases of infection. These signatures detected additional nodes at the XYZ site as well as many other client sites around the world.

To identify the level of ZeroAccess infection across its client base, NTT Group's security analysts and engineers developed a clone program to simulate behavior of the ZeroAccess botnet on an infected computer. When populated with 200 public IP addresses of known ZeroAccess nodes, the program identified over 360,000 worldwide supernodes. This information was used to further protect global NTT Group's clients.

Basic controls would have prevented the compromise of unpatched versions of applications and utilities on end-user systems.

Root Cause

It was apparent during review of the infected systems that the attacker had used a phishing campaign, enticing multiple XYZ employees to browse websites which were under the control of the attacker. Threats like ZeroAccess typically target end-user computers, not servers, though the initial infected machine in this instance was a server. Unpatched software on servers and end-user systems facilitated the initial compromise in the unmonitored networks, as well as the spread of malware across the internal network.

Cost of Incident

XYZ provided the actual cost of this event. NTT Group can project the expected cost of this event if XYZ had implemented other basic security controls.

Incident Management Profile	Cost of Incident
The actual cost of investigation, legal support, public relations support, remediation and professional incident support as described.	**$9,717**
If XYZ had fully installed and maintained their anti-virus software and maintained an active patch management system, this incident could potentially have been avoided.	**$0**

Figure 20: Cost of Event - ZeroAccess Botnet

Case Study Summary—ZeroAccess Botnet

Malware such as botnets pose a significant threat to any environment, especially when spread across an environment with inconsistent internal controls. ZeroAccess activity generated an IDS alert allowing a security analyst to identify suspect communications and detect the infection. Security analysts and engineers developed new signatures and gathered additional intelligence about the ZeroAccess botnet, to identify the exact nature and extent of the infection. SOC security analysts and XYZ worked together to isolate infections and prioritize response activities, ensuring mitigation was conducted effectively and efficiently. This case study shows XYZ was able to control their incident costs because of the rapid detection and successful response.

This case still presented opportunities for improved preventive and detective measures. NTT Group and XYZ used the incident as an opportunity to identify potential weaknesses in the XYZ environment.

Threat Mitigation—ZeroAccess Botnet

In this section we provide guidance on basic and advanced security controls to help protect against botnets and related threats.

The infection originated with users who were using server-class systems to browse the Internet.

Basic Controls – ZeroAccess Botnet Mitigation

- **Keep anti-virus software up to date.** The ZeroAccess botnet spread in the unmonitored client environment, partly because the organization's anti-virus software was outdated, and not all systems were being automatically scanned. Organizations should keep in mind that while anti-virus software is not a complete solution, it can be a valuable part of a layered security approach.
- **Enforce an active patch management process.** Patch management activities at XYZ were not assigned a high priority. Some systems were not patched to current levels, making them more susceptible to attack. If XYZ had implemented a more rigorous patch management process, they could have reduced the number of vulnerabilities and possibly eliminated the vulnerabilities exploited in this attack. Endpoint systems and commercial applications must be a part of effective patch management – failure to patch applications like Microsoft Office® and Adobe Acrobat® has led to very successful malware phishing campaigns over the last few years.
- **Implement comprehensive monitoring.** The ZeroAccess botnet was able to spread undetected through parts of the network XYZ had decided were less critical and were mostly unmonitored. The botnet was detected when it began to communicate with portions of the network which had more rigorous monitoring controls.
- **Employ 24/7 monitoring of alerts.** XYZ's botnet activity was identified within minutes by a security analyst after manual inspection of what seemed to be a policy violation which was deemed as "not critical" (peer-to-peer activity). 24/7 continuous monitoring by security analysts helps identify attacks within minutes, rather than hours or days.
- **Harden end-point systems.** Do not distribute end-user systems that include unnecessary and outdated software. Disable excessive administrative privileges. Implement controls to ensure end users cannot install software without authorization. ZeroAccess partially spread by exploiting vulnerabilities in unnecessary software.
- **Train and educate users.** Users provide the number one opportunity to prevent many incidents, especially malware-related events. In this case, it appears the infection originated with users who were using server-class systems to browse the Internet. Organizations should conduct frequent security awareness sessions for users and administrators to help them recognize malware, phishing and other attacks, as well as to ensure that all users understand their responsibilities in the incident response process.

Advanced Controls – ZeroAccess Botnet Mitigation

- **Use experts for malware analysis and incident response.** Advanced analysis with limited internal resources is often not practical. An organization should have outside expertise available to analyze botnets or other sustained attacks. Security professionals familiar with combating major security threats can quickly help identify the origin and extent of infections, and be used to rapidly identify mitigation approaches, such as development of updated IDS signatures.
- **Maintain a whitelist of approved applications.** The NTT security analyst was able to identify suspicious activity when he noticed apparent peer-to-peer file sharing activity, but XYZ listed no such approved application in their inventory. By maintaining an approved software list, an organization can more readily identify unauthorized software which could pose a risk.
- **Deploy and monitor purpose-built network malware detection capabilities.** Malware detection tools can detect complex communication and behavioral characteristics which may otherwise be difficult to track. When paired with an active anti-virus tool, such products can dramatically increase the effectiveness of malware detection.
- **Take targeted actions against ZeroAccess and other botnets.** An organization can take the following actions to directly address risks associated with ZeroAccess and related botnets:
    - o Adopt a layered defense strategy which includes segregating the organization's internal network into security domains. By segregating their environment at both the network and security level, an organization can define targeted controls for more valuable systems.
    - o Include egress monitoring on internal segments (to identify improper intra-segment communications) and for the external network (to identify improper communications leaving the organizational network).
    - o Monitor endpoint DNS traffic communicating on port 53 and restrict DNS queries to known DNS servers.
    - o Monitor for hosts sending and receiving UDP traffic on port 16464 and block them using proxy capabilities.

Web Application Attacks

In 2013, web application attacks ranked as the fifth most common type of attack identified by NTT Group. *Case Study: Massive Data Exfiltration via SQL Injection* describes how a single unsanitized data input field resulted in a prolonged web application attack, costing the client nearly $200,000 in incident response activities and the public disclosure of their customer database. The most significant categories of attacks impacting web applications are shown in Figure 21. SQL injection, XSS, CSRF, and other web application attacks represent serious threats to organizations and their data.

The following case study describes how a single unsanitized field on a web-based application cost one company almost $200,000 in incident response expenses.

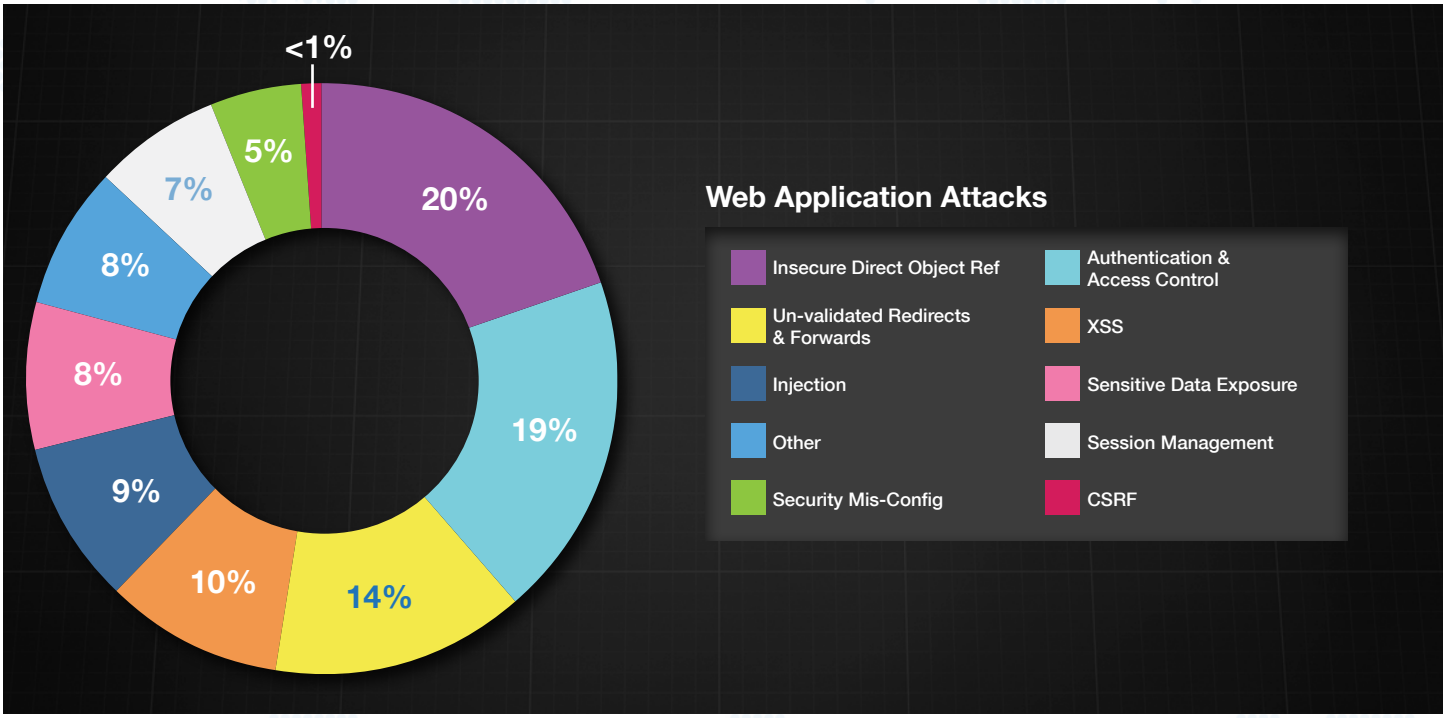

Figure 21: Types of Web Application Attacks

Case Study:
Massive Data Exfiltration via SQL Injection

CLIENT DOSSIER

Industry Vertical:	Financial
Location:	USA
Employee Count:	< 250
Financial Transactions (USD/Year):	$50 – 100 Million
Ownership:	Privately Held

ATTACKER DOSSIER

Location:	Eastern Europe
Primary Targets	U.S. Government U.S. Financial
Prefered Attack Methods	Cross-Site Scripting SQL Injection

CAPABILITY RATING

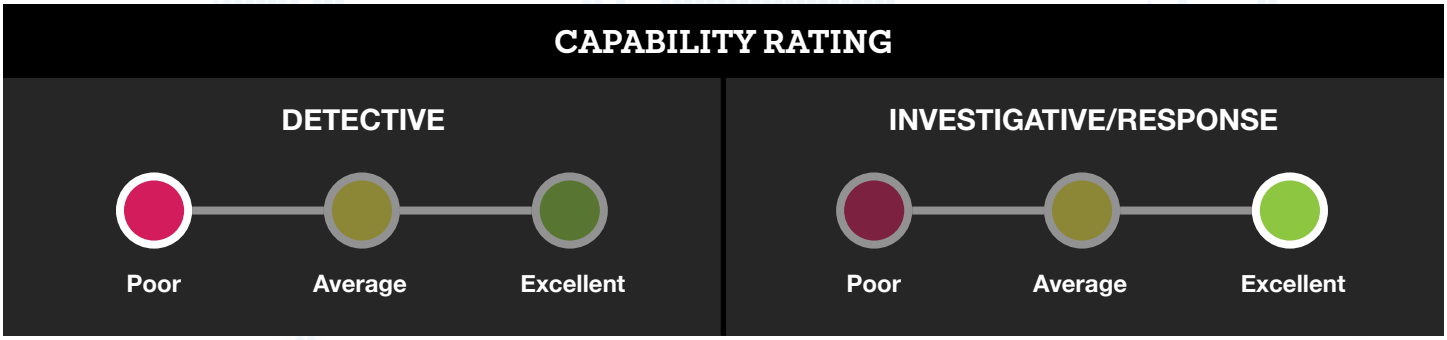

Case Study Moral:
Enforce basic controls in order to detect an attack.

Overview:
In November of 2013, XYZ National Bank (XNB) was notified by a third party that XNB client data was posted online. Investigation by NTT Group discovered that a prolonged SQL injection[4] attack had been underway for over 10 weeks. NTT Group assisted in rapid investigation and response to the attack.

[4] http://www.youtube.com/watch?v=71B2YNu9p6o

Timeline of Events

Date	Event Activity	Indicators
Day 0+	Initial site compromise	**Undetected:** Initial network reconnaissance activity, including port scanning and OS fingerprinting **Undetected:** Database reconnaissance activities, including table enumeration, mapping and data probing **Undetected:** Exfiltration of database informattion
Day 73	XNB informed by a third party that XNB client data had been posted to the Internet	**Detected:** Contact by a third party
Day 75	XNB hired NTT Group to perform forensic analysis of the event	
Day 76	Investigations confirm preliminary evidence of SQL injection attacks from five different IP addresses	**Detected:** Enumerating SQL statements from consistent outside sources identified as hostile by NTT Group IP address reputation database **Detected:** Previously unobserved, but logged activity from SQL injection, reconnaissance and exfiltration **Detected:** Verbose SQL injection activity following specific patterns
Day 81	XNB restores normal operations	**Detected:** Normal baseline traffic observed
Day 86	NTT Group asked to participate in a conference call with law enforcement concerning this incident	

Figure 22: Timeline of Events—SQL Injection

Description of Event

In November of 2013, XNB was notified by a third party that XNB client data had been posted online. XNB resources were unable to identify a breach and hired NTT Group to assist in investigation and response. XNB initially provided an incomplete set of logs from a subset of the involved systems. XNB web servers were not configured to capture full log data and thus could not identify the specific data exfiltrated.

Even basic monitoring of logs and alerts would have notified XNB at any time during the 10 weeks that they were actively under attack via SQL injection.

Forensic analysis of a second, more complete set of logs revealed the SQL injection attacks had been underway for as long as the logs had been retained (over 10 weeks). Even though early logs were incomplete, they showed indications of reconnaissance activity primarily centered on mapping of databases, tables, table fields and exploration of individual data records.

Forensic analysts determined the attacker had used an automated tool called Havij. Havij is an automated SQL injection tool which helps attackers perform automated SQL injection attacks, back-end database fingerprinting, database user and password hash retrieval, dump tables and columns, fetch data from the database and run SQL commands. It can even access the underlying file system and execute operating system commands.

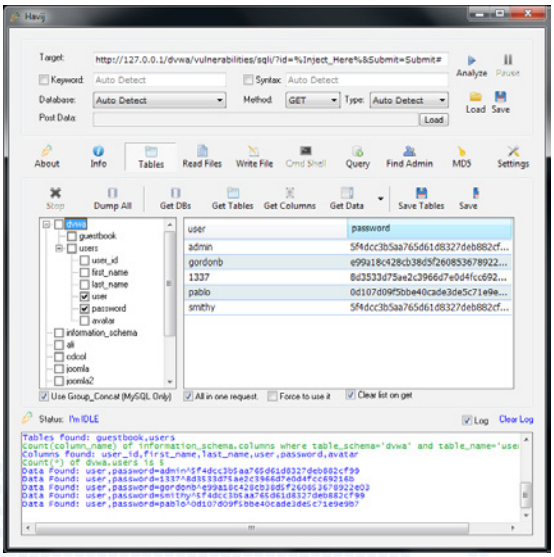

Figure 23: The Havij Tool

This type of attack automation and tool weaponization is a widespread trend. As a separate example, a web shell tool, used by hackers to remotely control compromised servers via web server interfaces, is shown in Figure 24.

Many of the tools used to perform automated SQL injection and other attacks are very sophisticated. Some resemble the complexity of legitimate, enterprise-class desktop and web applications software. Tools such as Havij not only have the capability to perform reconnaissance, but also allow attackers to escalate privileges, compromise additional systems, move laterally within the compromised environment, and exfiltrate data.

Figure 25 illustrates the capabilities of Havij with its built-in password cracking functionality, making this tool very convenient to use.

As with legitimate businesses, hackers invest in the tools used to increase productivity and automate repetitive tasks. Although we discuss some of the capabilities of Havij in this report, it is important to realize this is simply one of many tools with similar capabilities.

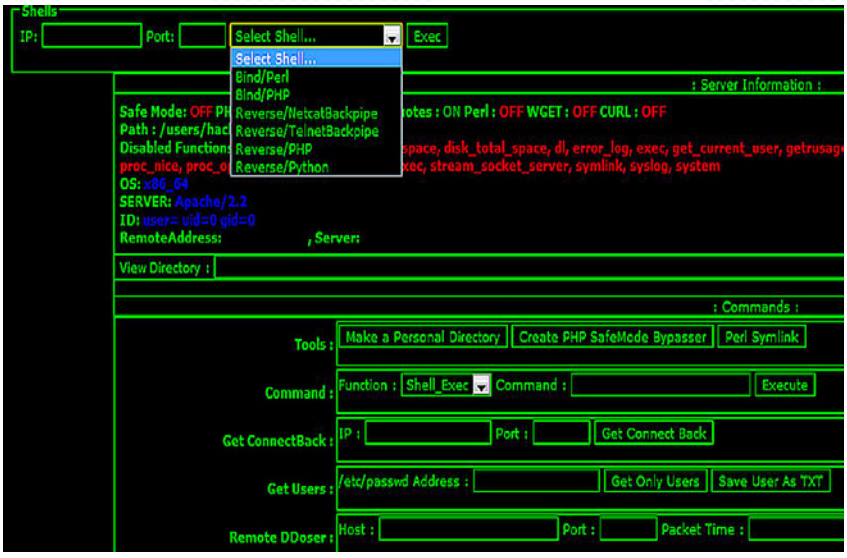

Figure 24: Web Shell Functionality

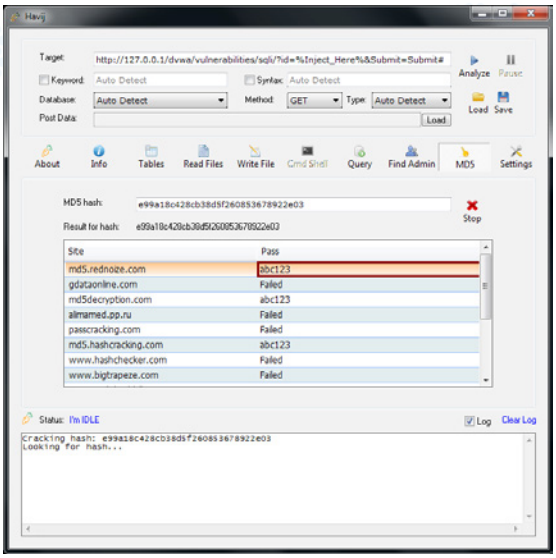

Figure 25: Cracking Passwords with Havij

Even basic monitoring of logs and alerts would have notified XNB during those 10 weeks that XNB was actively under attack. Advanced monitoring would have quickly identified the automated SQL injection, which would have enabled rapid, focused response. In the ideal case, no application or database would have been compromised and no data would have been exfiltrated.

Security analysts determined the majority of the data had been exfiltrated over a two day period at days 58 and 59 in the timeline in Figure 22. The attacker was able to access 52 tables across 11 databases, and extracted a large but unknown amount of data from those tables. Based on the amount of data available in the accessed tables and posted online, it appeared most, if not all, of XNB's client database had been uploaded. The attacker sent data to 14 different IP addresses, all of which were identified as hostile by NTT Group IP address reputation services. At the request of XNB, NTT Group passed investigation results to appropriate law enforcement agencies for further analysis.

Root Cause

Unfortunately, the breach occurred prior to the earliest date of any available logs, so analysts were unable to absolutely identify the root cause. However, the breached application still had a single field that was not being sanitized at the time of the engagement, and was still vulnerable to SQL injection. This field appeared to be the original point of entry. Incomplete logging and a complete lack of monitoring contributed to the sustained attack.

Cost of Incident

XNB provided the actual cost of this event. NTT Group can project the expected cost of this event if XNB had implemented other basic security controls.

Incident Management Profile	Cost of Incident
The actual cost of investigation, legal support, public relations support, remediation and professional incident support as described.	**$196,000**
If XNB had been doing reasonable monitoring of logs and detected the attack in its initial phases, they could have taken immediate and direct action.	**$24,980**
If developers had fixed the weakness that enabled the SQL injection attack at the time of development, or had supported the application with other protective techniques such as a web application firewall, the attack could have been completely avoided.	**$0**

Figure 26: Cost of Event - SQL Injection

Case Study Summary – Massive Data Exfiltration via SQL Injection

SQL injection is a common Internet-based attack used to compromise web-based applications. In this case, XNB was not prepared to either detect or investigate attacks. As a result, their time to detect the attack (executed over at least a 10 week period) and time to investigate the attack were both extended, resulting in an increased impact.

Threat Mitigation – Massive Data Exfiltration via SQL Injection

In this section, we provide guidance on basic and advanced security controls to help protect against SQL injection and other web application attacks.

Basic Controls – Massive Data Exfiltration via SQL Injection

- **Implement detailed logging for web applications and database transactions.** The attacker was able to conduct prolonged data gathering and extract a substantial amount of data, yet remained undetected by XNB. If XNB had utilized more robust logging, analysts could have more easily identified the attack in a timely manner. Better log data would have enabled higher quality analysis, supporting improved response.
- **Employ 24/7 monitoring of alerts.** XNB's compromise was identified by a third party who saw private XNB data posted online. Given how "noisy" this attack was, if XNB had been monitoring their environment in real-time they would likely have been able to identify the SQL attack immediately, preventing the compromise and data extraction.

- **Implement database access management logging.** The attacker was able to browse the XNB database and extract a substantial amount of information. Available logs could not clearly identify what data the attacker browsed and extracted. If XNB had implemented database access management they would have seen exactly what data had been compromised.
- **Conduct developer security training.** SQL injection is an application attack. In developer security training, developers learn to identify common design and programming mistakes which can lead to vulnerabilities. A properly constructed application would use parameterized queries and data validation to sanitize user input and reduce the likelihood of attacks such as SQL injection.

Advanced Controls – Massive Data Exfiltration via SQL Injection

- **Use experts for malware analysis and incident response.** Advanced analysis with limited internal resources is often not practical. An organization should have access to experts who can analyze an SQL injection attack or other sustained attacks. Advanced capabilities can more quickly identify the origin and extent of infections, and can also be used to rapidly identify mitigation approaches, such as development of updated IDS signatures.
- **Use web application firewalls.** A Web Application Firewall (WAF) will help detect and prevent attacks targeted against applications, such as SQL injection attacks. A WAF can even protect applications which have otherwise not been protected by the listed basic techniques. A WAF can support logging to help isolate the nature of such attacks.
- **Perform static and dynamic code review.** Code reviews allow developers to inspect execution paths and variables for potential misuse and increase the chances of eliminating errors which facilitate application attacks (including SQL injection attacks).
- **Perform application security assessments.** The XNB attacker compromised XNB through a vulnerability of which XNB was not even aware. If XNB had performed proper security testing of their environment, they could have identified the SQL injection vulnerability, and XNB would have had the opportunity to perform mitigation.

Core Security Management

This section provides an overview of core security management concepts and best practices. This guidance can assist organizations in addressing many of the fundamental security issues discussed in this report. We include the following core security areas:

- Minimizing the impact of threats
- VLM best practices
- Log monitoring best practices
- Incident response best practices

Minimizing the Impact of Threats

In this report, we go to great lengths to qualify and quantify the types of threats we see targeting our clients around the world by providing relevant case studies, statistics, charts, graphs and tools. We discuss advanced detective, investigative and response capabilities, as well as how emerging technologies like machine learning can be applied to solve security challenges. The biggest question we are asking is: **How do we minimize the impact of threats directed at our organization?**

We believe there are two basic strategies to minimize this impact: *avoidance* where possible and *response* when necessary. But there's a significant caveat to these simple concepts regarding the level of maturity, sophistication and capabilities of these strategies: while they may be simple in concept, they can be complex to execute.

We see organizations implementing basic security controls including anti-virus, vulnerability scanning, log monitoring, incident response, access controls, etc. What sets high-performing security programs apart is the manner in which those controls are selected, implemented and managed over the long term. Just because the controls may be viewed as basic doesn't imply their operation should be.

Threat Avoidance via Basic Security Controls

Basic security capabilities are widely available today. Many great products, platforms and standards exist to help organizations avoid threats. An underlying theme of many of these solutions is to "reduce the exploitable footprint" of your organization. This approach means looking at the risks and associated vulnerabilities in an environment and taking the actions needed to reduce risk. Some of these actions may include limiting access to network resources, hardening systems, patching, ensuring availability, logging events and planning for incidents.

How will doing the basics help your organization avoid threats?

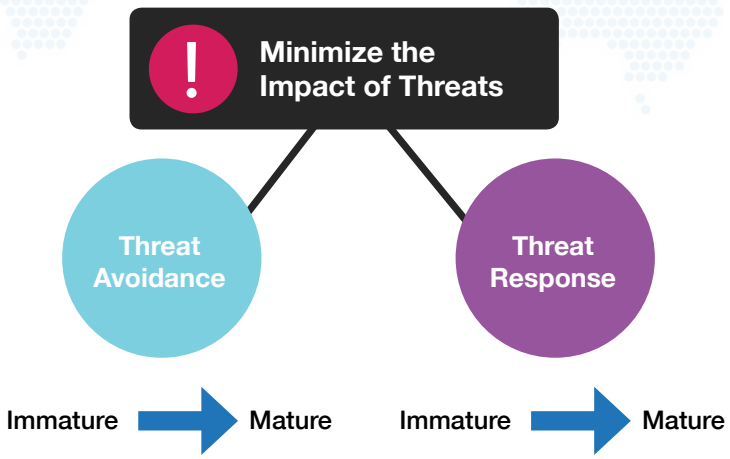

Figure 27: Minimizing Threat Impact

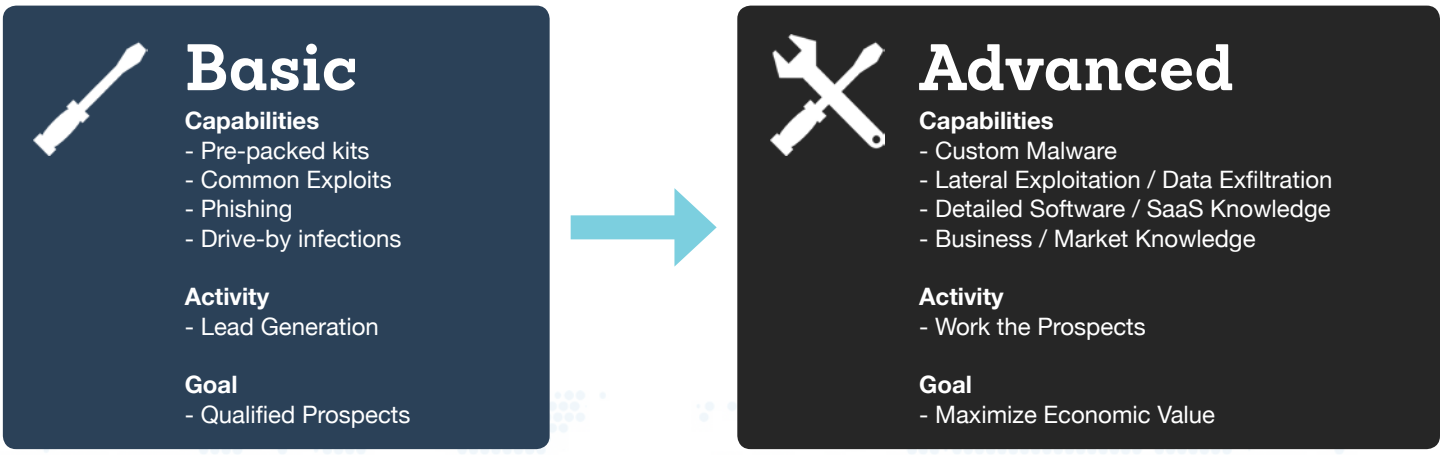

Figure 28: Basic vs. Advanced Capabilities

When we look at the vast majority of incidents impacting organizations, we see a pattern: malicious attackers, much like salespeople, perform "lead generation" using the broadest methods possible and basic attack tools to identify their most likely compromise "prospects." Only then do they begin to "work the deal" to extract the maximum economic value from the target.

Organizations doing a good job of threat avoidance are less likely to qualify as a prospect in the first place and have minimized the impact of the threat. Threat avoidance can be achieved using basic security controls, but these measures cannot simply be implemented and forgotten with the expectation that they will continue to work well.

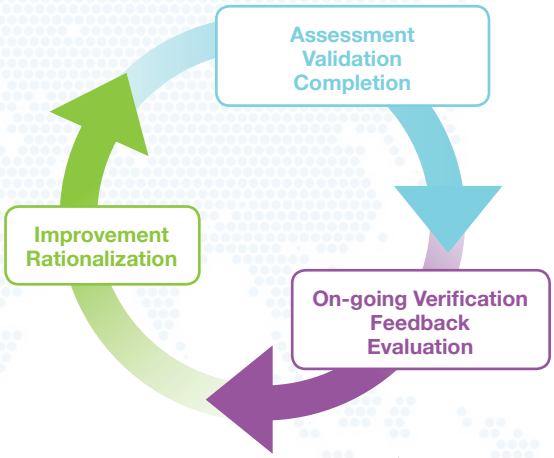

Figure 29: Lifecycle Management

An organization which selects, implements, and operates basic controls effectively will significantly reduce the number of threats by which the organization could be impacted. Done well, these controls build a strong foundation with an improved capability to respond to the threats that cannot be avoided and ultimately compress the mitigation timeline.

Doing the Basics vs. Doing the Basics Well

When implementing any security control, no matter how basic, ensure the control is going to work to the organization's maximum advantage. We recommend the following framework to guide the selection, implementation and operation (lifecycle management) of all security controls, both *basic* and *advanced*:

Assessment: Are you addressing the highest risks first? Have you performed a risk assessment and prioritized the risks?

Validation: Are you implementing the right controls? Many times the focus is on a selection between similar choices, when an entirely different type of control or a composite of multiple controls may accomplish the goals more effectively.

Completion: Has the control been completely implemented? Have you defined what "completely implemented" looks like? If not, is there an effort to ensure the scope of the control is improving over time?

Ongoing Verification: Is the control properly configured and functioning? You will often find a "set-and-forget" mentality which provides a false sense of security. Threats and regulations evolve and a static security program will not keep up with these evolutions.

Evaluation: Is the control effective? Are you testing it? Does it do what you envisioned it doing? Is it providing the value you expected? Are there potentially better alternatives?

Feedback: How do the results from this control relate to your other controls? What actions are you taking to do something with the results? Can you use the results more aggressively?

Improvement: How can you improve your use of the control? How can you improve your processes and procedures?

Rationalization: Can you simplify the number and complexity of controls without losing effectiveness? Can you reduce points of failure or increase robustness?

Your organization can apply this lifecycle management guidance to a security program which uses threat avoidance and threat response strategies with both *basic* and *advanced* capabilities.

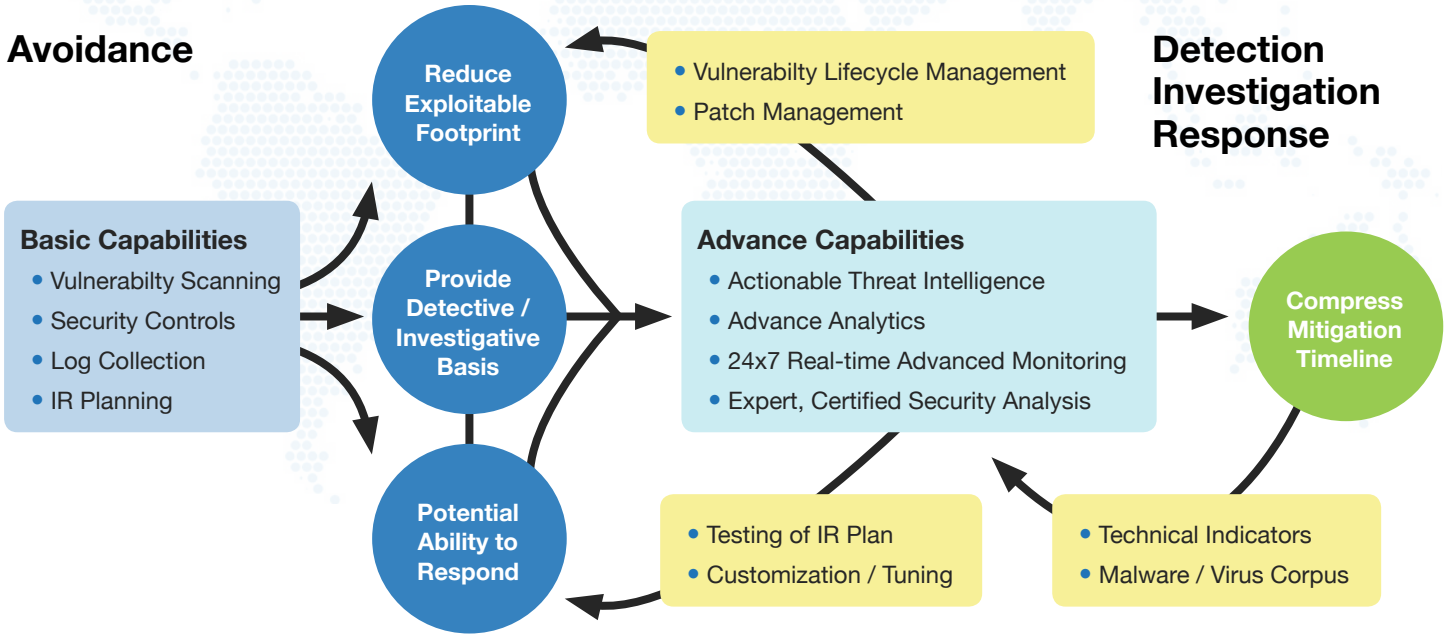

Figure 30: Basic and Advanced Outcomes

As Figure 30 illustrates, any security program should include a set of *basic* security capabilities. These *basic* capabilities should be planned, defined, implemented and operated in such a manner that they provide the key features of the security program including:

- *Reduce the Exploitable Footprint* by serving as security controls.
- *Provide Detective/Investigative Basis* by gathering intelligence about current operations.
- *Enable Ability to Respond* by serving as a foundational control for complementary functions.

It is easy to think about "advanced capabilities" from Figure 30 as technical hardware and software solutions; however, people, policy, process and procedure all play as much of a role, if not more, in implementing both basic and advanced capabilities. Advanced controls function best if they are built with a solid foundation of basic controls. These advanced security capabilities enable ongoing security management by supporting:

- Active security measures that allow improved detection, investigation and response, and compression of the mitigation timeline.
- Active security management functions such as VLM and ongoing patch management that also enable basic functions, allowing the entire security program to function more effectively.
- Active feedback on the effectiveness of the security program, such as information for tuning and results from testing of the incident response plan.

The detective, investigative and response capabilities of an organization directly rely on the threat avoidance capabilities in place. Those controls not only allow the organization to avoid certain threats altogether, but also ensure the underlying information needed during a security incident is available.

Reduce the Exploitable Footprint: VLM Best Practices

Evidence of Room for Improvement

An organization can proactively reduce the opportunity for an attacker to compromise their network. However, as identified in Figure 31, most organizations do a very poor job of using available intelligence to close vulnerable points in their infrastructure. 50% of vulnerabilities identified during scans against NTT Group clients in 2013 were first identified and assigned CVE numbers between 2004 and 2011. That is, half of the exploitable vulnerabilities we identified have been publicly known for at least two years, yet they remain open for an attacker to find and exploit.

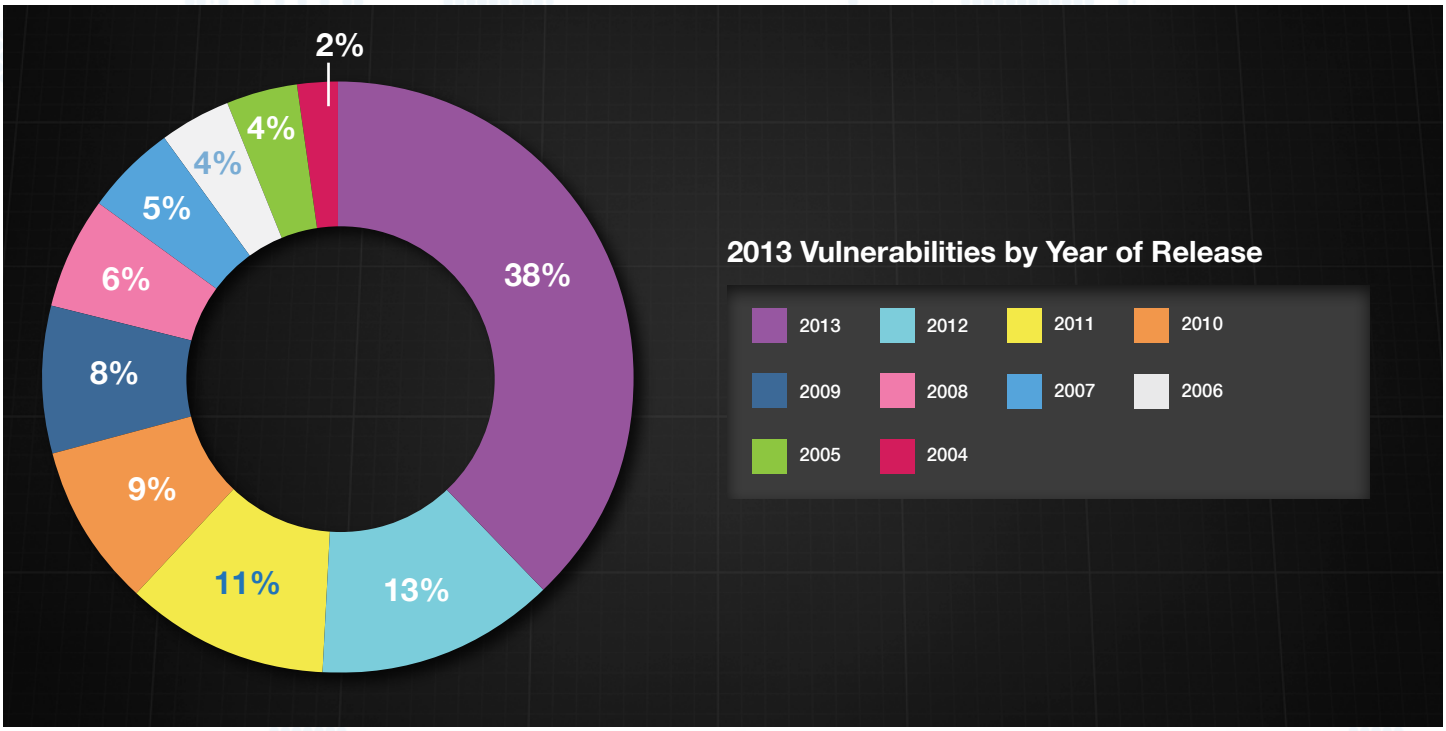

Figure 31: Vulnerabilities by Year of Release

This data indicates many organizations today are unaware, lack the capability, or don't perceive the importance of addressing these vulnerabilities in a timely manner. Maturity in patch and configuration management is the best way to approach this type of situation.

PCI Has Had a Positive Impact on Threat Avoidance, However…

PCI requirements and controls are encouraging organizations to secure their external perimeter. Organizations which perform quarterly external PCI Authorized Scanning Vendor (ASV) assessments show fewer vulnerabilities, as well as much faster remediation times. As seen in Figure 32, among organizations using a mature VLM process, those with PCI compliance requirements have a 27% faster remediation time than organizations which do not. This number increases to 35% when compared to those not using a VLM process at all.

But as we will discuss later in this section,organizations with PCI requirements aren't necessarily "better" at addressing vulnerabilities, even though they are demonstrably faster.

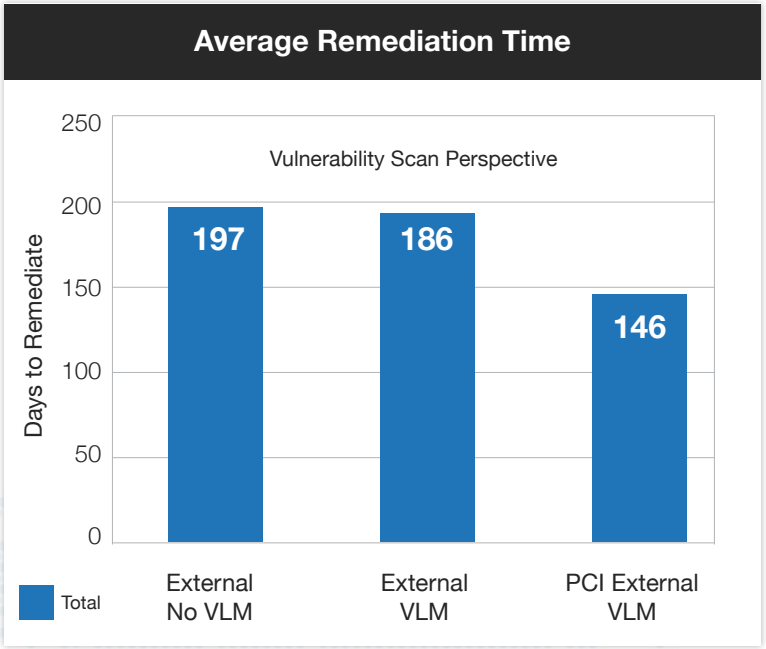

Figure 32: Average Remediation Time

The Case for Effective Vulnerability Lifecycle Management

Organizations using vulnerability management capabilities not only have fewer vulnerabilities, but also reduce their average time to mitigate those vulnerabilities. In reality, we still see two types of organizations:

- Those which take a scan-and-dump approach to vulnerability scanning and tend to haphazardly remediate vulnerabilities
- Those which regularly review, manage and set a disposition for their vulnerability data (false positive, remediated, risk reduced, risk accepted, etc.)

The immaturity of security programs in scan-and-dump organizations results in higher numbers of vulnerabilities and longer mitigation timeframes.

As seen in Figure 33, organizations with active vulnerability management programs are far less exposed to the threat of exploit kits. Those with mature VLM programs accounted for only about 20% of all instances where vulnerability scan data directly correlated with the capabilities of exploit kits.

Many organizations still have far to go in implementing a mature VLM program. The common process of *scan and dump* and using Microsoft Excel® to manage remediation efforts, is simply inefficient. A more formal method of tracking vulnerabilities, including use of a VLM system, can help an organization reduce its external and internal vulnerabilities and reduce its mitigation timeframe, resulting in a noticeably more secure organization.

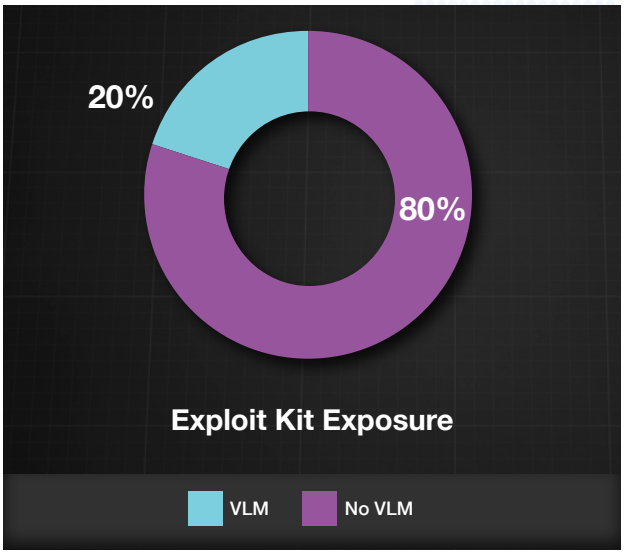

Figure 33: Exploit Kit Exposure

Four Common Approaches

NTT Group sees organizations prioritize their VLM process using some blend of the following methods:

- **Risk-based.** Many organizations take a simplified risk-based approach, reviewing the vulnerability scan results by severity, starting at the highest severity and working their way down.

  A simple risk-based approach can be inefficient, and potentially dangerous, because it treats all the highest severity vulnerabilities as equal. In reality, they are not equal, and this reinforces a "find it, fix this one, move on" mentality which ignores the root cause. An organization must account for the scope of systems affected as well as the root cause of the vulnerability. Making one small change to an OS configuration which addresses 20 outstanding high and medium vulnerabilities in a short time has a much higher ROI than spending days addressing a vulnerability that affects a single system. A better risk-based approach is the "find it, understand root cause, fix it forever and for everything, move on" approach. Organizations with this mentality spend less effort addressing vulnerabilities and are more likely to avoid future threats.

- **Asset-based.** Some organizations take an asset-based approach. These organizations consider which vulnerabilities are detected on critical assets / subnets and remediate those first. This is typically blended with the risk-based approach, which is good, but root cause analysis still typically takes a back seat to the "fix it now" pressures the security organization faces.

- **Operational.** Organizations that take an operational approach to VLM analyze the data by fix type (patches, OS / application configuration, etc.) then build a "to-do" list for the appropriate operational team and assign ownership to the team. This process includes plans for the organization's next vulnerability re-scan to measure how the operational team is doing.

  This approach is often seen in large organizations with a high degree of separation of duties and many disparate operational teams. It often lacks the prioritization of either the risk-based or asset-based approaches, and misses the root cause analysis entirely. This usually results in mitigating the "easiest to quickly address" vulnerabilities, regardless of severity.

- **PCI-based.** PCI recommends organizations take a risk-based approach towards remediating vulnerabilities. These organizations often take a strict pass/fail approach – if the vulnerability would cause a failing report, that's what the organization fixes first. For example, there is an issue with a Denial of Service vulnerability assigned a CVSS score of 10.0, but considered "passing" by PCI, and is virtually ignored by most organizations because it wouldn't affect their PCI compliance status. The PCI Security Standards Council strengthened QSA testing procedures to ensure organizations are validating that these results are factored into the security program, but this practice still remains an area of weakness for most organizations.

A small percentage of organizations take a more modeled approach using security device/network management tools. This methodology can provide a better risk-modeled approach to mitigate vulnerabilities, but is still not a common practice. Likewise, we rarely see organizations employ a more technical risk-based approach by correlating which high-severity vulnerabilities should be fixed before other high-severity vulnerabilities. This level of analysis is based on threat information, including: common exploits, exploit kits which contain the vulnerability, hacking tools which target the vulnerability, and potential attack vectors.

Three Pitfalls To Avoid

Three pitfalls await organizations looking to improve their VLM process:

- **Reactive vs. Proactive.** Vulnerability scans should help validate the organization has accomplished actions to make them more secure, and not be viewed as a comprehensive means to achieve good security. Organizations tend to look at vulnerabilities as something that happened to them rather than understand *why* they happened to them. Obviously some vulnerabilities are the result of weaknesses in the firm's software, but in other cases the vulnerability results from actions taken (or not taken) by the organization.

  It is crucial an organization understands the root cause of vulnerabilities and identifies any security controls which failed or were missing. Is the patching program behind? Is the System Development Life Cycle (SDLC) security testing not identifying vulnerabilities as part of the development process? Are hardening guides insufficient? Vulnerability scans should confirm you have a well-structured security and operational process. Organizations which use vulnerability results to secure networks in a reactive mode are missing the substantial benefits which root cause analysis can provide.

- **Not taking advantage of available tools.** In the best of circumstances performing VLM well is difficult and requires ongoing diligence. Failing to use available tools to automate the risk scoring analysis, root cause analysis and the coordination and documentation of mitigation efforts makes a difficult job much harder. A key advantage of many tools is the ability to reduce the ongoing burden as vulnerabilities are addressed over time. If your organization's VLM is primarily a spreadsheet-driven activity requiring many man-hours, you are spending more resources to address fewer vulnerabilities. There are many commercial off-the-shelf, enterprise and SaaS tools and services which can significantly reduce operational costs and dramatically increase effectiveness.

- **Authenticated vs. "hacker's view".** Vulnerability scanning can be performed as either uncredentialed (the "hacker's view" where the scanner must perform all checks remotely) or authenticated (with credentials that provide access to the host). There are times when credentialed scans aren't feasible, such as on highly-secured environments. The base OS may be hardened to the point that the services needed for remote authentication and auditing aren't enabled. For everything else, an organization should use credentialed scanning, but they commonly do not. Some organizations have a focus on getting a "hacker's view" of their vulnerabilities. Theoretically, this enables focused effort, but the approach has significant drawbacks:

  - Credentialed scans drastically reduce the number of false positives within vulnerability results. Reducing the false positives means less time chasing vulnerabilities which don't exist and more time fixing actual ones.
  - Many scan vendors have vulnerability checks and verifications that can only be performed using credentials for the target system. Why not get the most value from your scanning tools?
  - Many scan vendors can audit and validate end-point configuration and settings, inventory software and enforce best-practice standards (e.g., NIST, CIS, Microsoft, USGCB) when credentials are provided. When used as part of management for basic security controls, these results can assist the ongoing verification function.
  - Hackers have plenty of help from end-users to get into your organization. Focusing only on what's vulnerable from the outside ignores the reality that most attacks involve an initial foothold and then lateral exploitation to maximize the economic benefit for the attacker. This "Maginot Line" mentality has been shown throughout history to have significant weaknesses.

Periodically Review and Refine Your VLM Strategy

VLM shouldn't just be about reporting the number of critical, serious, or informational vulnerabilities in January and repeating the same statistics in February. It should be an analysis of what vendors, products, security controls or processes caused the vulnerabilities and what was done to address them. The best organizations emphasize what can be done to prevent similar vulnerabilities in the future.

VLM results are only meaningful if they actually help you manage the vulnerabilities, and not just count statistics. Not all vulnerabilities pose an equal threat to the organization, and not all vulnerabilities can be closed with the same level of effort. If sheer numbers are important it may be because upper management does not understand the amount of time and effort remediation requires.

Conclusion: VLM is Not Just About the Vulnerabilities

VLM shouldn't be about "find it, fix it and move on." That common approach misses the opportunity to avoid future threats through increasing the baseline security of deployed systems. VLM should be about understanding what the root cause of the vulnerability was, addressing that cause, remediating the vulnerability and validating the fix across repeated scans (ongoing verification of the security control).

Provide Detective/Investigative Basis: Log Monitoring Best Practices

Evidence of Room for Improvement

As part of the on-boarding process for log monitoring Managed Security Services (MSS), organizations are required to provide details of devices, platforms, applications and databases to be included in the log monitoring scope. Based on experience with this process, about 75% of organizations need to perform (or be assisted with) significant discovery during this process.

75%

75% of organizations need to perform (or be assisted with) significant discovery during this process.

Organizations which have a primary objective of meeting compliance requirements tend to have a more accurate assessment of what is in scope for their environments.

During and after discovery, approximately 50% of organizations realize the initial contracted scope is not sufficient for several reasons:

- The security team discovers devices, platforms, applications and databases that were previously unknown to them.
- Miscommunication exists between the management objectives in performing log monitoring and the technical team's understanding of how to meet those objectives.
- The original scope does not fully meet compliance requirements.

Applications and databases pose particular challenges during the on-boarding process due to the requirement to bridge the gap between the security team and the application and database teams.

Finally, one third of the organizations have at least one of their devices, platforms, applications or databases configured in a manner that does not provide the security information required to meet their needs.

The Case for Effective Log Monitoring

The first notification an attacker is taking advantage of an organization's vulnerabilities can usually be found in their device logs, if those systems are capturing the appropriate log information.

System and application logs are among the most valuable, yet underutilized, security resources available. Generated logs are often too overwhelming for manual review, may not be set up to log the necessary information, or can be intimidating to set up effectively. Where compliance standards require periodic review or log retention for a defined time, logs may be viewed as simply a checkbox requirement.

In many forensic response cases, investigators find the evidence needed to identify the malicious activity cannot easily be found in the logs of the affected systems, if those logs even exist. An example of how insufficient logging and log management helped contribute to a prolonged attack is included in *Case Study: Massive Data Exfiltration via SQL Injection* (a basic attack that continued undetected for more than 10 weeks).

Managed properly, system and application logs can be a rich source of security information, providing administrators with a detailed outline of events and actions. Logs can tell a compelling story—who is accessing devices, what changes are being

made, what the device is seeing on the network, what is happening within the environment. Logs may show activity on other systems interacting with the device generating the logs. Logging provides detailed insight into anomalies and threats occurring in your environment.

Logs do not show the organization what someone thinks is happening; they show what is really happening. This is illustrated in *Case Study: Administrator Releases a Worm*, where if the organization had been paying attention to their logs, they could have seen their remediation actions had been unsuccessful, and changed their approach instead of continuing their failed attempts for almost eleven weeks.

While the content and usefulness of logs varies widely, getting the most security value out of your logs requires the actions described in the following sections.

Enable and Monitor Logs on the Devices Providing the Information You Need Most

All logs are not created equal. The devices your organization chooses to monitor, and what events to log, depends on many factors – industry, regulatory requirements, previously encountered security issues, concern over specific threats, value and vulnerability of data or systems. Organizations must focus on collecting and monitoring the logs which provide the most value for their own environment. If the organization's primary concern is web application attacks such as SQL injection and cross-site scripting, firewall logs will be of limited value. Such an organization is better served by focusing on logs from web servers, web applications firewalls, databases, middleware and intrusion detection systems.

Figure 34 shows how different types of devices and log settings may be used in different areas of your environment to create multiple layers of logging and monitoring. A layered security approach provides even more value when considering the complementary logs an organization can get from successive layers of logging.

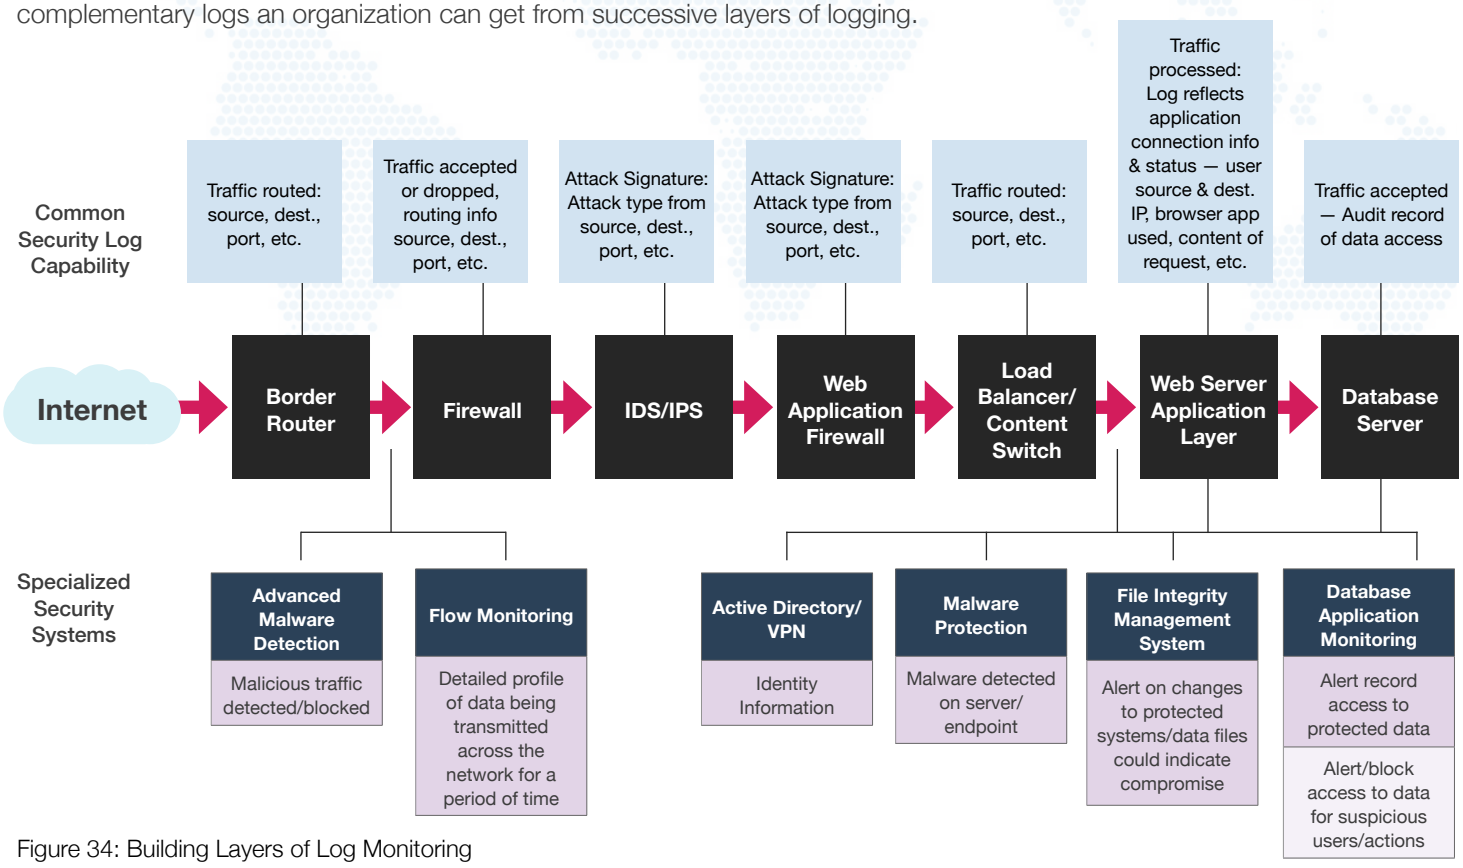

Figure 34: Building Layers of Log Monitoring

Your organization should use log data to paint a picture of what is happening in your environment.

Configure Log Settings to Enable Adequate Logging

The organization must first identify their most important logs and log sources. It can then configure the type and level of logging that provides the appropriate level of detail. A good rule of thumb is to set up logging that can provide the four Ws: Who, What, When and Where.

At the same time, be careful not to enable logging for data that is not useful. On most perimeter firewalls, logging millions of accepted connections each day is not only excessive, but may also overload the device and cause logs to be lost. Log sizes from databases, servers (AS400 series servers, for instance, can generate very high volumes) and devices can vary dramatically based on the implemented configuration settings, so organizations must identify information to be captured, and capture what is needed. Most devices allow for several levels of logging, either via different criticality levels or customizable fields within the logs. Defaulting to the highest level of logging is easy, but can upset the delicate balance between capturing the logs you need and being overwhelmed with enormous volumes of information.

The need to configure appropriate logging can be demonstrated by the sheer number of logs processed by NTT Group on behalf of clients. NTT Group security companies monitor over 200,000 devices for thousands of clients. Figure 35 shows the average daily log counts generated by the top ten log generating technologies.

There are clear differences between the volume of logs generated by different technologies. If a typical database generates over two million logs daily, organizations need to ensure that they are capturing only the data they want, have mechanisms to sort through the logs to identify events of interest, and have archiving capacity to meet compliance and investigative requirements.

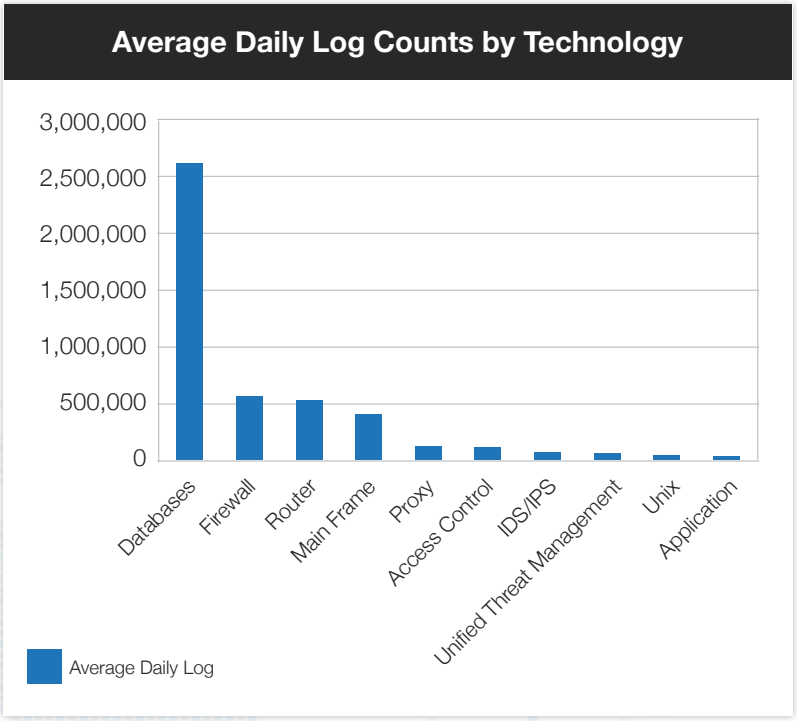

Figure 35: Average Daily Log Counts by Technology

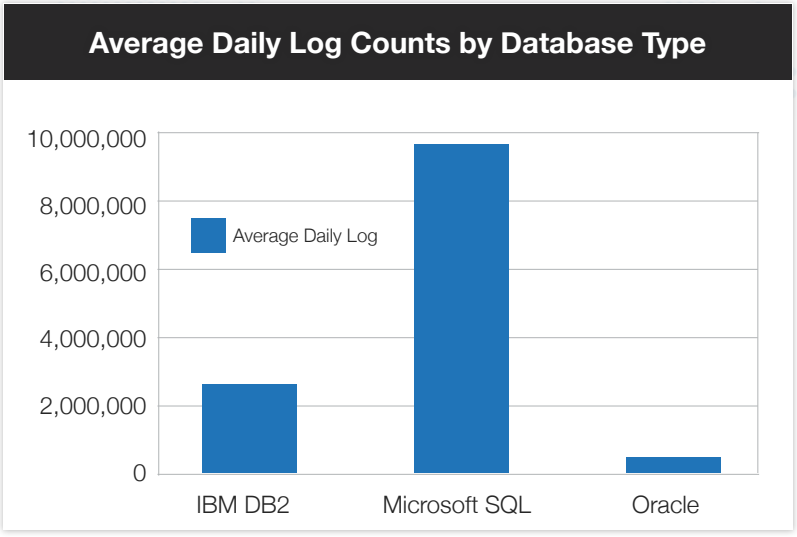

Figure 36: Average Daily Log Counts by Database Type

Even within technologies, there can be dramatic differences in the number of logs generated.

In environments we monitor, Microsoft SQL server databases generate almost 10 million logs daily, while Oracle databases generate less than 200,000 logs daily. By default, a Microsoft SQL Server generates a significant number of logs. While Oracle provides granular audit settings, these are not typically enabled.

Compliance standards such as PCI, SOX, HIPAA/HITECH and others may define specific information which must be logged and the amount of time it must be saved.

Review, Normalize and Tune Log Data to Assure Quality of Information

Logs are sometimes overly chatty, redundant or meaningless without added context. It is important to normalize log data, remove false positives and tune data to gain the best possible insight. Such logging support is a basic, foundational security control. It is even more valuable to turn normalization and tuning into a regular, repeatable process to ensure the data is always actionable when needed. Lifecycle management of the log monitoring system can turn a basic IT tool into a valuable security asset.

While it is important that an organization perform initial normalization and tuning, it is even more valuable to turn this into a regular, repeatable process that will ensure the data is always actionable when needed.

Normalization requires an organization to collect logs for an appropriate time period before they can perform an adequate review. To determine the appropriate period for log normalization, consider several factors:

- The function of each device and the area it covers. For example, employee tasks and shift hours may cause log volume fluctuations which should be taken into account.
- Daily, weekly, monthly or other periodic tasks and applications could affect the data to be normalized. A nightly backup may cause an increase in log traffic which might look alarming if not understood in context.
- Scheduled assessments, maintenance windows and average volumes of reconnaissance traffic. These occasional events and outliers can result in logging anomalies as well, and, if possible, should be considered when normalizing the data.

False positives are log lines that may appear to show there is a problem, but are related to normal activity. For example, logs might indicate a server is experiencing a brute force password-guessing attack, when the actual issue is that a user's password has expired. Dealing with false positives is important, but care must be taken not to suppress other, actionable log data while doing so.

An organization may need to filter false positives out, or fine tune individual signatures. This tuning step is an often-overlooked aspect of normalizing data. While it may be good to reduce data volume, the data can still be meaningful for incident response or forensic analysis.

Initial normalization is essential to creating a baseline and working toward actionable alerts, but the work isn't finished there. New and evolving threats, vulnerabilities, detection signatures and anomaly patterns require constant attention to ensure logs remain useful and your environment is adequately protected. It is also important to account for network changes, organizational moves and vendor or application swaps to ensure the baseline remains current.

Transmit Logs To a Central Repository for Storage, Analysis, Reporting and Alerting

Small, quiet environments may have few logs generated, which could be reviewed manually and stored on the system that generates the logs. However, in most situations logs are generated in volumes which cause difficulty for review, storage/retention processes, compliance and forensic purposes.

To address this difficulty, many organizations choose a centralized log solution. This may be in the form of a network-based log storage location, log management server, security information and event management (SIEM) system, outsourced SIEM management or a managed security services provider (MSSP).

Log centralization offers many advantages, especially the management of storage. On a firewall or server, logs may be overwritten in weeks or in seconds, depending on the logging policies and storage space available. Once overwritten, they are lost forever, along with the valuable system information they once held. A scalable storage solution on a centralized device allows logs to be retained in larger volumes and for a longer duration. Compression and encryption can add to the value of this centralized log storage. Secured central storage can help meet specific compliance requirements; for example, section 10.5.4 of the PCI-DSS 3.0 standard calls for writing logs from external-facing technologies to centralized internal storage, helping to prevent loss or alteration.

Another advantage of log centralization is the ability to use the aggregated logs from various devices for analysis, cross-correlation, reporting and in some cases real-time alerting. Consolidated log information can feed reports for management and compliance purposes. And while logs may display signs of trouble on an individual system, the aggregated, cross-correlated logs from multiple systems can identify threats which might not be visible in a single system. Using centralized log analysis, suspicious activity in a server's event logs might be correlated with firewall logs showing the download of executables from a known malware site, confirming a likely infection and allowing rapid response to the intrusion. The centralized solution should be capable of alerting key personnel when threats are recognized.

Use an Effective Log Transport Mechanism

When setting up a centralized log management solution, an organization needs to plan for the secure transmission of those logs, which will often contain sensitive data.

The most common log format for non-Microsoft devices is syslog. Many legacy devices which support syslog will transmit logs using UDP. Unfortunately, UDP is a connectionless protocol which provides no error recovery, no guarantee of delivery and sends log data in clear text. Depending on the device or Log Transport Agent (LTA) used, you may be able to select TCP instead. TCP is a connection-based protocol which helps to ensure delivery and provides the ability to buffer and retransmit data when necessary. Some newer devices and applications also support SSL or TLS encapsulation for secure, encrypted transmission.

While TCP is more reliable over the network than UDP, both are still susceptible to loss in the application's buffer. I/O contention can lead to data being dropped from a buffer or queue before the system has a chance to transmit its log data over the network.

Enrich Log Data to Make It Even More Useful

While some logs do a good job of capturing exactly what occurred, others can leave an administrator searching for answers. To increase the value of log information, some solutions will enrich log data through capabilities such as network awareness, Active Directory integration and cloud-based reputation services. Log enrichment is a great start to help make the leap from log monitoring to alert handling.

Consider the issues associated with monitoring in a DHCP environment with thousands of users. When logs indicate a problem, an IP address might suggest which device to look at, but if the analyst knows the associated domain IDs, it would facilitate communication with affected users. This additional correlation and context can reduce the amount of manual effort required and allow analysts to focus on more critical tasks.

Asset information can associate log data to an individual device, the device's function, geographic location, regulatory requirements and the responsible department. If vulnerability information is available, this can be correlated to validate whether a system may be vulnerable to certain attacks.

Properly linking the data (cross-device correlation) can add context *inheritance* between devices. Some systems and applications have a high degree of context (typically users, applications and security devices) because they retain more information about their environment. Some devices have a low degree of context (typically basic network devices like routers and switches) because they only understand the data they see. Without proper correlation, this context may not be available for all logs.

Periodically Review and Refine Your Logging Strategy

Environments change. New threats emerge. Logging and monitoring strategies need to change as well, taking into account new systems added or obsolete systems that have been removed, applications that have changed, and new or updated compliance standards. Make a habit of reviewing your lifecycle management (see "Doing the Basics vs. Doing the Basics Well" in this report for more guidance on lifecycle management of security controls).

Conclusion: Logging is Not Just About the Logs

To make the best use of log data, an organization must:

- **Identify the correct scope.** Ensure logs are captured from appropriate devices in the environment.
- **Identify correct settings.** Ensure the granularity captured in logs strikes the correct balance between capturing all required information and collecting unnecessary data that makes pertinent information more difficult to identify, wasting storage and bandwidth.
- **Consolidate and correlate logs.** In some instances, elements selected in logs are meaningful on their own; in other cases, correlation is necessary to bridge logs from different devices, enriching the collected information. Consolidation and correlation of logs from different devices into a single view provides a more comprehensive picture of the environment.
- **Enrich data.** Log data enrichment with user information, geographic, or IP reputation information can provide important context for analysis.

Logging Example Detail

Figure 37 shows how a combination of logs from disparate systems can tell a coherent story. In this sequence of logs, a device is infected by malware and begins communicating with a command and control server. By consolidating and correlating logs from four different log sources, a complete view of activity from the infected system is possible.

Log ID 1 and 2 - 9:52:48 – File integrity monitoring (FIM) software and anti-virus protection simultaneously log activity. FIM software logs the modification of a protected file, while anti-virus detects the malware. Both alerts originate from IP address 192.168.10.100.

Log ID 3 and 4- 9:52:48 – Simultaneously, alerts are received from both a network IDS and a Cisco Security Appliance that IP 192.168.10.100 is attempting to access a blacklisted site associated with previously identified malware activity.

Log ID 5 – 9:53:49 – The network IDS identifies traffic associated with botnet C&C activity.

Log ID 6 – 9:54:49 - Alarms are triggered from firewall logs, based on the volume of traffic being generated from the system at 192.168.10.100.

Log ID	Time	Summary	Dest	Port	Source	User Name	Log Source
1	09:52:48	2010281 — File Modified			192.168.10.100	Rsullivan	OSSEC File Integrity Monitoring
2	09:52:48	2012481 — Virus Found by Symantec SEP			192.168.10.100	SYSTEM	Symantec Endpoint Protection
3	09:52:48	2024061 — Blacklisted Activity Detected	xx.xx.233.2	80	192.168.10.100	Rsullivan	Snort/Sourcefire IDS
4	09:52:48	2024061 — Blacklisted Activity Detected	xx.xx.233.2	80	192.168.10.100	Rsullivan	Cisco Adaptive Security Appliance
5	09:53:49	16743 — Bot C&C Traffic	xx.xx.253.152	80	192.168.10.100	Rsullivan	Snort/Sourcefire IDS
6	09:54:49	2018461 — Outbound Traffic Exceeded Threshold	xx.xx.253.151	80	192.168.10.100	Rsullivan	Cisco Adaptive Security Appliance

Figure 37: Logging Example

Enabling the Ability to Respond: Incident Response Best Practices

Evidence of Room for Improvement

During 2013, NTT Group security teams responded to many client incidents. As shown in Figure 38, 77% of the organizations involved had no incident response team, policies or procedures in place to effectively respond to a significant cyber incident.

Furthermore, of the 23% of organizations which had some incident response planning in place, very few were mature or well managed. Unfortunately, this number reflects the ad hoc nature of many incident response plans. Incident response is a fundamental security control and should be updated on a consistent basis.

> The less an organization speculates about what they think is happening to their network the better. It is always better to rely on hard facts and evidence.

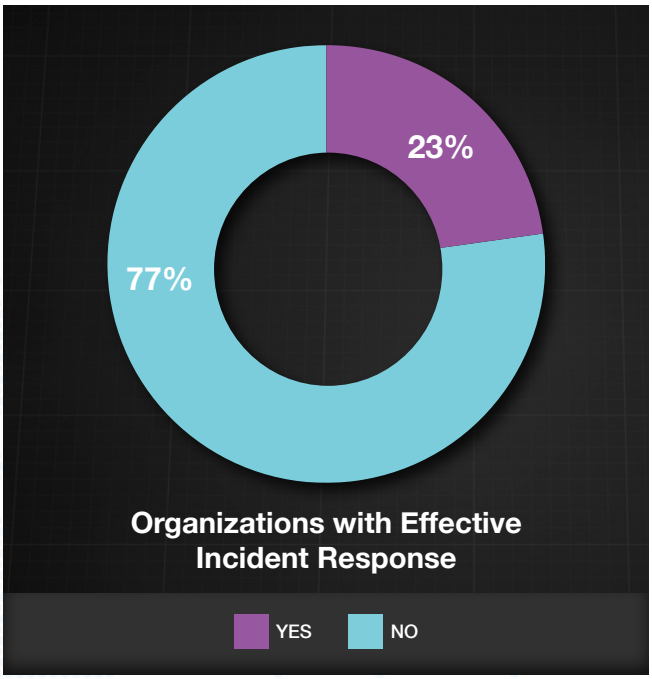

Figure 38: Organizations with Effective Incident Response

The Case for Effective Incident Response

Regardless of the exact type of attack or its source, the effectiveness of an organization's incident response plan can dramatically affect the impact an incident will have on an organization. *Case Study: Administrator Releases a Worm* describes how an organization's lack of effective incident response planning cost $109,000 in incident response efforts plus months of continued infection and degraded service. NTT Group received frequent client requests for incident response support during 2013. Almost all of the client requests were of an urgent nature, relating to an ongoing attack and represented the first interaction between the client and NTT Group Security response teams. As shown in Figure 39, 43% of the engagements were related to suspicious activity later attributed to malware, or direct knowledge of an active infection within a host or network segment. Many of these were common malware variants which could have been avoided if effective basic prevention methods had been in place. Advanced methods could have quickly identified those infections which had bypassed basic controls.

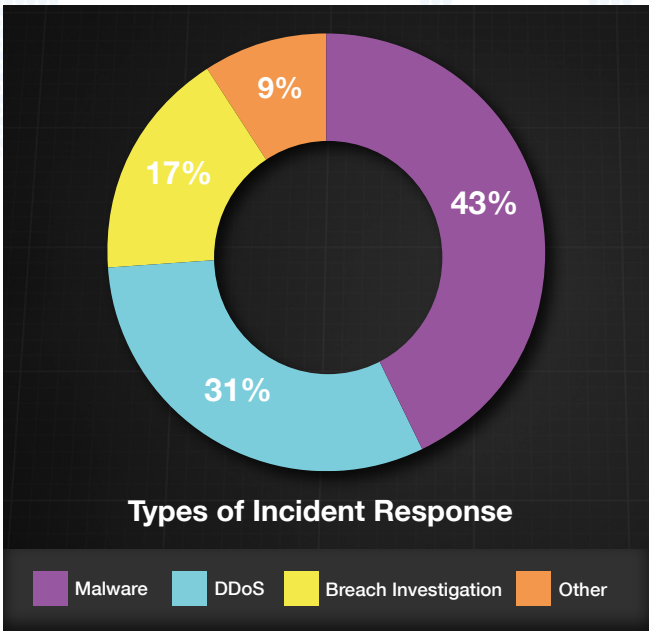

Figure 39: Types of Incident Response

DDoS mitigation and follow-up investigation represented 31% of incident responses. This statistic indicates prevention methods are not adequately implemented and many organizations are not equipped to handle an active DDoS attack. A targeted network breach, which is the type of attack most people associate with a security incident, represented only 17% of the investigations in 2013.

Incident Response = Detect + Investigate + Respond

When responding to an active threat, organizations have three areas in which they can implement controls to reduce potential loss:

- Detect
- Investigate
- Respond

Most organizations do not have budget defined for handling a significant cybersecurity incident. As a result, many organizations spend precious response time attempting to secure emergency funds.

Note: This information assumes organizations have already implemented basic security controls such as firewalls, proxies, purpose-built controls and other general principles of effective defense. Every organization's environment is different, so recommendations are included for organizations to consider and evaluate.

An organization must detect an incident before it knows what to investigate. Effective investigation tells an organization what they have to know about the incident in order to conduct a meaningful response. If any part of the process fails, the entire response process is likely to fail. This section includes examples of *controls observed and which may be implemented by any organization, but does not represent a complete list of all possible controls or processes.*

Detective Capabilities

The ability to identify an attack is a fundamental requirement of being able to investigate and respond. The less an organization speculates about what they think is happening to their network, the better. It is always better to rely on hard facts and evidence. *Case Study: Massive Data Exfiltration via SQL Injection* highlights the impact an undetected incident can have on an organization. It underscores that basic log management is fundamental to any security program and that active log monitoring can improve detection of active threats. *Case Study: ZeroAccess Botnet* describes how full function, active monitoring might have helped an organization identify a system compromise before they became an active ZeroAccess supernode.

Any information your organization can obtain about potential attacks will give you a tactical advantage in dealing with a potential incident. An investment in threat intelligence offerings such as IP reputation, brand monitoring, emerging threat notifications and vertical market-focused threats can significantly augment your understanding of current threats.

There is no shortage of services and products to enable visibility into an organization's environment. The controls and technologies in Figure 40 include some of the services and products an organization should consider.

Investigative Capabilities

Investigative controls take the information identified from incident detection and verify the organization has the information required to conduct a meaningful response, including:

- Validating the detection of an incident is accurate and not a false positive
- Determining the scope of the incident, including the assets affected and the degree to which they are affected
- Conducting advanced analysis of incident indicators to determine the appropriate level of response
- Validating the action plan is effective as remediation is implemented

Response Capabilities

Response capabilities determine how effectively an organization can reduce the impact and minimize loss potential of a confirmed incident. Response capabilities usually involve many stakeholders and technologies that must function together. It is vital for an organization to plan in advance and communicate effectively.

Part of this planning is ensuring support vendor relationships are predefined. This definition includes understanding the scope and understanding of ISP, MSSP, anti-virus capabilities and obligations to support the organization during an incident. An organization should review and test incident response plans at least annually and should proactively include third-party vendors. After real-world incidents, or even training exercises, an organization should conduct a formal post-incident review to document lessons learned and address gaps in the plan. *Case Study: Administrator Releases a Worm* highlights how an ineffective incident response cost one organization over $109,000 in response activities and resulted in almost four months' worth of degradation of service, troubleshooting and associated mitigation effort.

Logging Infrastructure

- IDS, IPS, WAF, Proxy, Web, Database, Firewall
- Authentication, Directory Services
- Load balancer, Critical system logs
- MSSP Provider
- *Logging Best Practices* section of GTIR

Log Monitoring

- Add context to logs (assets, geo, business purpose, etc.)
- Correlate across devices, business units
- Active analysis, reporting and tracking

Intrusion Detection/Protection System

- Provides network visibility for known attack signature
- Allows rapid signature development and deployment
- Flexable and mature defense control
- Capability to dynamically mitigate attacks (IPS)

Network-Based Malware Detection

- Purpose built for malware threats
- Email, network and content security capabilities
- Often have forensic analysis capabilites

Threat Intelligence Gathering

- Vertical market focused and general feeds widely available
- IP reputation
- Brand monitoring
- Identify emerging threats
- "Tell me its going to happen, before it happens"

Budgeting for an incident is no small task. It is often hard to convince organizations to invest for an incident which may not happen. NTT Group security companies' experience indicates that most organizations do not have budget defined for handling a significant cybersecurity incident. As a result, many organizations spend precious response time attempting to secure emergency funds and fast-tracking purchasing, contracts and legal issues.

Compressing the Mitigation Timeline

The best opportunity to compress the mitigation timeline, providing the best chance to reduce the impact of threats, comes from combining threat avoidance and threat response capabilities into a strategic approach. These capabilities must be built on effective basic and advanced security controls that are appropriate for the organization.

An organization's detection, investigation and response capabilities can be mapped directly to the three most intensive phases of formal incident response. These capabilities directly impact the organization's ability to reduce loss.

In the following figure, we illustrate the three phases in a linear fashion. First, we detect the attack. Then we move to investigation, and based on the outcome of the investigation we conduct our response (if we validate the incident is legitimate *and* worthy of a response). In this discussion, we refer to this as the *Overall Mitigation Timeline*.

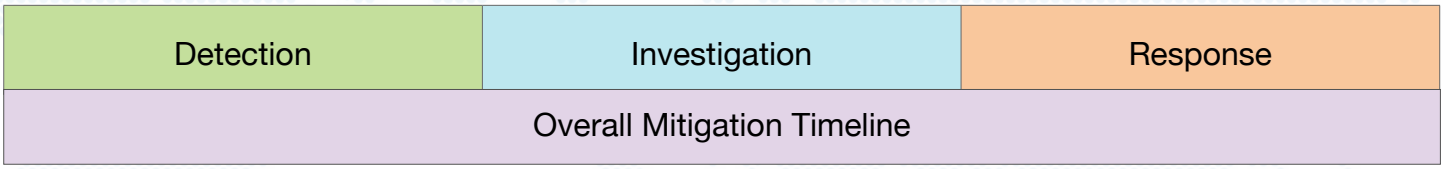

Figure 41. Overall Mitigation Timeline

The Detection, Investigation and Response sections in Figure 41 are the same length. As an example of how your organization's actions and capabilities can affect the duration of the mitigation timeline and resulting loss exposure, we will use the example of a common incident experienced by many organizations during 2013, a DDoS[5] attack. The data in Figure 42 is derived from a real world case study found in the 2013 Global Threat Intelligence Report.

2013 Global Threat Intelligence Report – DDoS Case Study Data	
Source IP Addresses	**91,435**
Countries Involved	**150**
Elapsed time before attack detected	**2.5 hours**
Investigation time	**0.5 hour**
Response time	**10.5 hours**
Approximate loss	**$67,500**

Figure 42: DDoS Data

[5] *http://www.youtube.com/watch?v=2rhf0FqOoLQ*

Figure 43 shows where the organization spent most of its time and money during mitigation. At a calculated loss of approximately $5,000 per hour, and a total loss of $67,500, we can see the organization incurred most of the cost during their detection and response efforts.

Figure 43. Poor Detection - Poor Response

To consider the overall mitigation timeline, we have to understand the reasons behind these numbers.

Q: Why did it take 2.5 hours for the DDoS to be detected by the client?

A: The client had very poor detection capabilities for the segment of the network being targeted by DDoS. The client had detection capabilities in place for its PCI in-scope environment but neglected to apply any such capabilities to its corporate website. In this case, the client did not realize its website was under attack until they started receiving a high volume of calls from clients who could not access the site.

Q: Why did the response effort last 10.5 hours?

A: After the incident passed the investigation phase, the client initiated defensive activities by filtering some of the DDoS traffic; however, their defensive capabilities were unable to keep up with the high rate of network packets from the DDoS attack. The client contacted their ISP to implement upstream filtering of the attack traffic. Unfortunately, it took another 5.5 hours to reach appropriate personnel at the ISP who could implement the filtering.

Q: What if the client were able to reduce the amount of time for one or all of the phases?

A: If an organization can reduce the length of any phase in the incident response process they can potentially compress the mitigation timeline. This is best seen in an example of how being efficient, and having the proper capabilities, could significantly reduce the impact of this incident.

Let's assume the client had previously built out a proper detection capability, and had been able to detect the DDoS in a fraction of the time from our previous example. Instead of taking 2.5 hours to detect the attack, the organization is now able to detect the DDoS attack within 15 minutes of the start of the attack. Figure 44 shows how this reduces the detection time by two hours and 15 minutes, and reduces the total loss to $56,250 ($11,250 less than previously experienced).

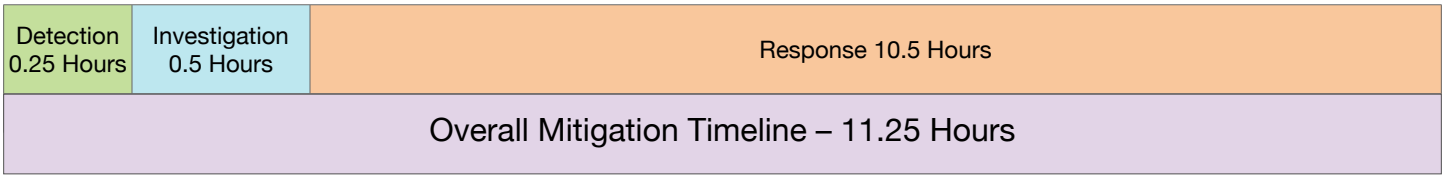

Figure 44. Good Detection - Poor Response

If the client had not changed any of their detection capabilities, but instead, by investing in reducing the response time alone, had been able to conduct its response in 2.5 hours, this would save eight hours of response time. With this detail shown in Figure 45, they would reduce the loss to $27,500 ($40,000 less than the initial incident).

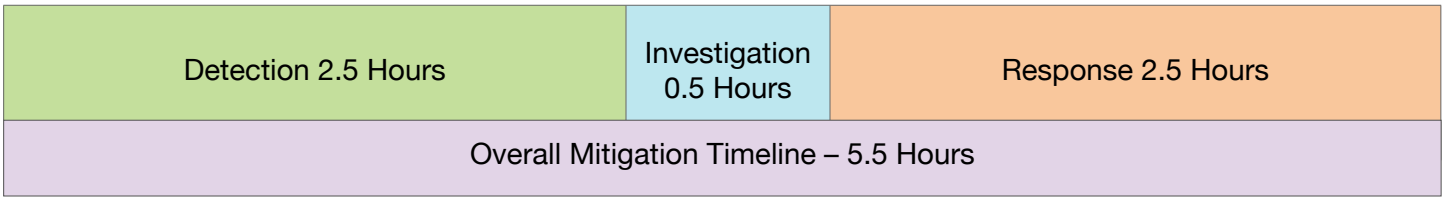

Figure 45. Poor Detection - Good Response

The recommended approach is to invest in reducing detection, investigation and response times to ensure all three capabilities work together for rapid and effective mitigation. In Figure 46 we see that the overall length of the attack is now 3.25 hours (by saving 2.25 hours of detection time and eight hours of response time).

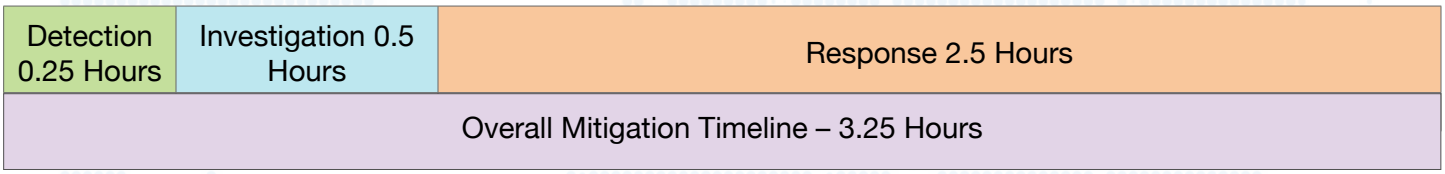

Figure 46. Good Detection - Good Response

If implemented prior to the attack, these capabilities would have resulted in a loss of only $16,250 ($51,250 less than the original incident.) This demonstrates it is wise to ensure your mitigation plan is well-defined, clearly communicated and tested in advance of an attack.

To put things into clearer perspective: consider that in many cases, incidents such as these do not occur once in a year, but many times. Consider the impact that reducing the mitigation timeline would have on this client if they were attacked 10 times in a year instead of once (total incident cost of $162,500 instead of $675,000).

For this reason, it is best to look at risks using a formal risk assessment methodology. With a proper risk assessment, your organization should have a very good idea of what risks pose the greatest threat, have the highest probability of occurrence and could have the largest financial impact.

Consider carefully the difference between acceptable risk during the normal course of business and acceptable risk while under attack. During normal business, a security control that could keep a handful of clients from accessing your systems may be deemed unacceptable; however, in the face of an ongoing attack your organization may look differently at the situation. What wasn't acceptable during the normal course of business now becomes the lesser of two evils and the difference between apologizing to some clients, or all of them.

The risk assessment process will help your organization decide where to invest its security budget. No one will claim a risk assessment is easy; it is often a humbling experience if conducted correctly, as it will truly illustrate your loss potential, based on quantitative and qualitative risk analysis.

Certainly, risk cannot be 100% removed. Organizations will usually develop mitigating controls to minimize impact, but still allow for services to be rendered to clients. There is always risk associated with business goals, and knowing what risks to accept, reject, defer or mitigate are key parts of doing business.

No one will claim a risk assessment is easy; it is often a humbling experience.

There are many risk assessment methodologies to choose from. NTT Group recommends performing some research before adopting your approach or having a vendor perform a risk assessment service for your organization. Regardless of your risk assessment mechanism, the results of the assessment can help define the detective, investigative and response controls that will provide the best chance to compress your mitigation timeline.

Next-Generation Detection

Applying Machine Learning for Real-Time Monitoring

In the recent past, we have seen continued growth in malware-based cyberattacks designed to evade signature-based security controls. Attackers are employing sophisticated techniques to ensure their malware distributions can avoid detection, enabling the malware to spread throughout an organization and cause widespread compromise of critical systems. Threats which are difficult to detect using anti-virus software or security appliances (e.g., IPS, WAF, SIEM) are increasing at a rapid rate.

NTT Labratories research shows that current anti-virus fails to detect 54% of new malware collected by honeypots and 71% of new malware analyzed by sandboxes.

NTT Laboratories' research shows current anti-virus fails to detect 54% of new malware collected by honeypots and 71% of new malware analyzed in our sandbox environment. Of course, most malware would be detected within the next few days as signatures get updated with new threat information; however, some malware is still not detected for a full 10 days or more, if ever. This delay can be attributed to the fact anti-virus software vendors require time to detect new types of malware and distribute updated signature files. During this delay period, users are at risk until they are able to employ enhanced signature files, and by then, the malware may have already exposed the organization to additional threats.

This situation can clearly be improved by enhanced real-time detection which we believe can be enabled through the application of machine learning capabilities.

Why Machine Learning?

Current malware signature patterns are very dynamic. New malware is released regularly, old malware is re-released and repurposed, and tools to update malware signatures are common, requiring anti-virus vendors to rapidly develop and distribute new detection signatures. Unfortunately, methods currently employed are relatively static, relying on identification of known signatures which are periodically distributed in database updates. This process impairs the detection of malware whose signature pattern changes even slightly. This limitation prevents signature-based malware detection from keeping up with new malicious attacks.

Machine learning capabilities, on the other hand, introduce the ability to detect new attacks based upon the trained recognition of malicious behavior patterns. For example, we can consider features to recognize communication type, volume and content. This type of information can be used to detect malware behavior in a predictive manner, something which cannot be accomplished through any current signature-based malware system.

Historically, processing capabilities have limited the success of machine learning capabilities. This meant machine learning could only be applied to a controlled amount of data and could only focus on a defined set of characteristics within the data. This limitation has prevented the application of machine learning capabilities in the MSSP environment, given the enormous amount of log data to be analyzed and the number of characteristics within the data.

In the past, the processing power required to apply machine learning capabilities to high-volume MSSP environments was a technological challenge. NTT i^3 applies Jubatus, a machine learning framework from NTT Laboratories, to achieve real-time analysis and training capability. Jubatus has already been used in advanced malware analysis and the machine learning engine has proven itself capable of identifying patterns and relationships which would have been very difficult, if not impossible, for a human analyst to identify. As these capabilities develop, they will be applied to real-time monitoring and detection of new threats against the MSSP landscapes.

Future Plan

NTT i³ is engaged in a number of initiatives focused on enhancing machine learning capabilities in the threat monitoring and detection environment. Although we recognize machine learning requires a high-quality training phase, we have already observed positive results as we continue with this endeavor.

A machine learning engine (Jubatus) has proven itself capable of identifying patterns and relationships which would have been very difficult, if not impossible, for a human analyst to identify.

References

[1] http://jubat.us/en/

[2] http://www.ffri.jp/assets/files/research/research_papers/psj13-murakami_EN.pdf

About/Client Contact

Solutionary, an NTT Group security company (NYSE: NTT), is the leading pure-play managed security services provider (MSSP), focused on delivering managed security services and global threat intelligence. Comprehensive Solutionary security monitoring and security device management services protect traditional and virtual IT infrastructures, cloud environments and mobile data. Solutionary clients are able to optimize current security programs, make informed security decisions, achieve regulatory compliance and reduce costs. The patented, cloud-based ActiveGuard® MSSP platform uses multiple detection technologies and advanced analytics to protect against advanced threats. The Solutionary Security Engineering Research Team (SERT) researches the global threat landscape, providing actionable threat intelligence, enhanced threat detection and mitigating controls. Experienced, certified Solutionary security experts act as an extension of clients' internal teams, providing industry-leading client service to global enterprise and mid-market clients in a wide range of industries, including financial services, health care, retail and government. Services are delivered 24/7 through multiple state-of-the-art Security Operations Centers (SOCs).

See how Solutionary can enhance security, improve efficiency and ease compliance. Contact an authorized Solutionary partner or Solutionary directly at 866-333-2133 | info@solutionary.com | www.solutionary.com

NTT Com Security, an NTT Group security company (NYSE: NTT), is in the business of information security and risk management. By choosing our WideAngle consulting, managed security and technology services, our clients are free to focus on business opportunities while we focus on managing risk.

The breadth of our Governance, Risk and Compliance (GRC) engagements, innovative managed security services and pragmatic technology implementations, means we can share a unique perspective with our clients – helping them to prioritize projects and drive standards. We want to give the right objective advice every time.

Our global approach is designed to drive out cost and complexity – recognizing the growing value of information security and risk management as a differentiator in high-performing businesses. Innovative and independent, NTT Com Security (formerly Integralis) has offices spanning the Americas, Europe and APAC (Asia Pacific) and is part of the NTT Communications Group, owned by NTT (Nippon Telegraph and Telephone Corporation), one of the largest telecommunications companies in the world.

To learn more about NTT Com Security and our unique WideAngle services for information security and risk management, please speak to your account representative or visit: www.nttcomsecurity.com/us/ for regional contact information.

Dimension Data, an NTT Group security company (NYSE: NTT), and one of the world's largest security systems integrators, delivers broad technical and integration expertise across a variety of IT disciplines, including: networking, security, communications and collaboration, data centers, virtualization and end-user computing. With over 14,000 employees and operations in 52 countries across five continents, we manage more than USD12.5 billion of network infrastructure through five Global Service Centers on a 24/7 basis, in more than 15 languages. We service over 6,000 security clients across all industry sectors, including financial services, telecommunications, health care, manufacturing, government and education. Our real-time security information and event management architecture is based on an enterprise-wide risk management solution that enables our SOC analysts to centrally manage attacks, threats and exposures by correlating security information from multiple security technology controls. This solution enables our analysts to eliminate clutter such as false positives, while quickly identifying the real security threats to help them respond effectively and efficiently. Our team of certified security experts,

located in Security Operations Centers (SOCs) across five continents, brings unmatched cybersecurity experience to augment the knowledge base of our clients' IT organizations. We provide peace of mind with skilled technicians ready to help clients respond to, and mitigate, all cybersecurity threats. Our certifications include ISO9001, ISO 27001, ASD Protected Gateway, PCI DSS, Cisco MSCP, ACSI 33 and ASIO T4.

For more information, please contact your nearest Dimension Data office or visit www.dimensiondata.com.

NTT Innovation Institute, Inc., (NTT i³) is the Silicon Valley-based innovation and applied research and development center of NTT Group. The institute works closely with NTT operating companies and their clients around the world to develop market-driven, client-focused solutions and services. NTT i³ builds on the vast intellectual capital base of NTT Group, which invests more than $3.5 billion a year in R&D. NTT i³ and its world-class scientists and engineers partner with prominent technology companies and start-ups to deliver market-leading solutions that span strategy, business applications, data and infrastructure on a global scale. To learn more about NTT i³, please visit us at www.ntti3.com.