The Next Generation IPS

Comprehensive Defense Against Advanced Persistent Threats

Contents

Brought to you compliments of:

Introduction

Today's enterprises and government agencies are faced with advanced persistent threats (APTs). These attacks often use social engineering and other techniques to gain a beachhead on corporate networks. From there they probe and use "privilege escalation" to gain access to high-value information and then use stealthy techniques to send the information outside without detection. The attackers are far more patient than the perpetrators of yesterday's simple smash-and-grab operations. Often these attacks are too sophisticated to be countered by today's typical information security defenses.

But a new form of defense, the Next Generation IPS, combines application awareness, context awareness, content awareness and high-speed intrusion prevention capabilities to detect and stop many APTs.

This white paper will examine the following:

- What are advanced persistent threats?
- How APTs work: The RSA breach of 2011
- Defining a Next Generation IPS
- A three-layer defense model
- How a Next Generation IPS might have stopped the RSA breach
- Protecting Against APTs: The HP Enterprise Security Network Defense System

What Are Advanced Persistent Threats?

Enterprises and government agencies are increasingly menaced by new varieties of APTs associated with organized cybercriminals and state-sponsored hacker groups. These attackers have far larger budgets and longer time horizons than the small-time vandals and con artists responsible for earlier generations of threats such as worms and distributed denial-of-service attacks. They are responding to vastly bigger incentives: multimillion-dollar financial gains and corporate and government espionage.

Today's organized cybercriminals have figured out how to take advantage of new technologies, trends and attack techniques, including:

- Social media that makes it much easier for criminals to gather personal information they can use to launch highly personalized social engineering attacks.
- Mobile devices that can be used to gain footholds in corporate networks.
- Zero-day attack techniques that provide ways to compromise laptops and PCs before threat signatures and security patches can be deployed.
- Narrowly targeted attacks that fly below the radar of antivirus vendors, so signatures and patches are never created at all.

Unfortunately, these trends make it increasingly easy for new attacks to evade conventional information security defenses. Social engineering techniques dupe employees into giving attackers passwords and planting malware on corporate systems. Most antivirus packages, firewalls and first-generation intrusion prevention systems (IPS) rely on threat signatures and patches that are difficult to disseminate to all endpoints (including mobile devices) and that may never be developed at all for targeted attacks.[i] Signature-based technologies are still a necessary part of layered defenses, but they will never be able to cover all of the attacks launched today.

How APTs Work: The RSA Breach of 2011

In March 2011, RSA, the security division of EMC Corporation, publically disclosed that a sophisticated cyberattack had penetrated its internal networks and extracted information from its systems related to its SecurID two-factor authentication products. The consequences for RSA were serious, including $66 million in direct costs to replace customers' security tokens, monitor customers and harden its infrastructure. The breach was also a blow to RSA's reputation.[ii]

To its credit, RSA published a number of details about the attack that allow us to see how a multi-part APT operates — and how it might have been stopped sooner. Let's examine the attack, step by step.[iii]

Penetration phase

The attack on RSA unfolded in three phases. In the initial penetration phase, social engineering techniques were used to plant a malware payload on a system on the internal network.

1. The attacker sent phishing emails with the subject line "2011 Recruitment Plan" to four employees at EMC, the parent company of RSA (Figure 1). The email appeared to be from the webmaster of Beyond.com, a career network web site.

Figure 1: The phishing email (Source: F-Secure)

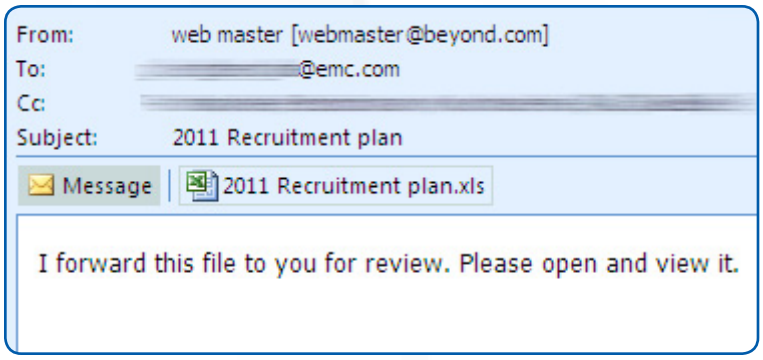

2. One of the employees retrieved the email from his junk mail folder, read the two-sentence message ("I forward this file to you for review. Please open and view it") and opened the attached Microsoft Excel file "2011 Recruitment plan.xls." The Excel file was blank except for a small symbol in one cell (Figure 2).

Figure 2: The spreadsheet with the embedded Flash object (Source: F-Secure)

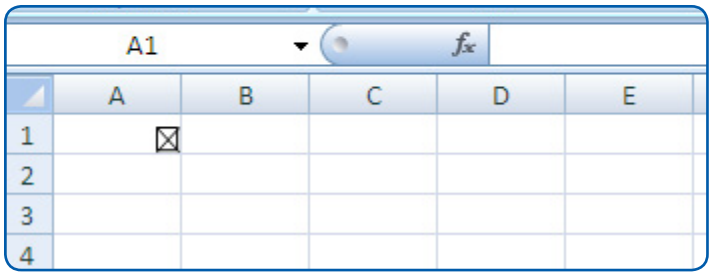

3. When the spreadsheet was opened, an embedded flash object exploited an Adobe Flash vulnerability (CVE-2011-0609) to plant a remote access Trojan (RAT) called Poison Ivy on the employee's system.

Exploitation phase

In the exploitation phase of the attack, the attacker:

4. Used the Trojan to establish a command-and-control link from the comprised internal system to an external system.
5. Probed and mapped the network to find strategic systems containing target information.
6. Found users with high-level administrative privileges and captured their credentials.
7. Used the credentials to gain access to systems with the target information (a process known as privilege escalation, as it involves achieving higher levels of access privileges).

Exfiltration phase

In the final exfiltration phase, the attacker:

8. Extracted confidential data from the compromised high-value systems (while also launching diversionary attacks on the company's HR systems to distract RSA's security team) and moved the data to servers that could be used to stage the data for external transmission.
9. Compressed and encrypted the data into RAR files on the staging servers, then used FTP to transfer the files to an external system (good.mincesur.com) that was widely used for cyberattacks.
10. Retrieved the information from the external system and used details about RSA's authentication technology to develop attacks on RSA customers.

Let us make two observations about this breach:

- There were several points in this multi-step attack where, *in theory*, the damage *could* have been prevented.
- Human nature and technical weaknesses meant that *in practice* conventional security defenses *did not* thwart the determined, patient attacker.

It is very likely that RSA had deployed a full complement of antivirus packages, firewalls and first-generation intrusion prevention systems. Almost certainly RSA's employees and IT staff were at least as security-conscious as their peers in other companies. Yet the company suffered a painful and very public breach.

And this type of incident is by no means unique to RSA. RSA began to track down the penetration in less than two weeks. Many victims don't detect the attacks for months, and sometimes not until they are alerted by outside third parties. At one large telecommunications company, hackers remained active for *six years*.

So, is there a defense against this type of clever, persistent attack? Can attacks like this be detected and stopped?

Defining a Next Generation IPS

Garter analysts John Pescatore and Greg Young address the issue of defending against advanced targeted threats in a research note titled, "Defining Next-Generation Network Intrusion Prevention."[iv]

Pescatore and Young start by noting, "Advanced targeted threats are using evasion techniques and new delivery methods that are penetrating existing network security defenses." Although they do not mention the RSA breach, it is clear that they have that type of attack in mind.

Pescatore and Young then enumerate five characteristics required in a Next Generation Network IPS that can cope with the new threats:

- **Standard first-generation IPS capabilities** such as being able to read network traffic at wire speed and block attacks that can be recognized on the basis of threat signatures.
- **Application awareness and full-stack visibility** to identify traffic from specific applications and enforce network security policy at the application layer.
- **Context awareness** to use information from a variety of sources outside of the IPS to improve blocking decisions — for example, user identity information from directories, geo-location information such as the source of packets, and reputation feeds identifying suspicious web sites and IP addresses.
- **Content awareness** to inspect and classify executables and other file types in inbound and outbound traffic.
- An **agile engine** that makes it easy to upgrade to new information feeds and detection techniques developed in the future.

Pescatore and Young see a major need for systems with these characteristics. They predict that "the success of advanced targeted threats in evading existing IPS products and causing significant enterprise damage will drive rapid demand for better IPS capabilities."

A Three-Layer Defense Model

Pescatore and Young's requirements include a new set of capabilities that extend intrusion prevention systems in new directions. These systems must provide much higher performance, the ability to scan more layers of the network stack (such as the application layer) and the ability to handle many more types of context- and content-related rules.

But their prescription implies a solution that goes beyond the central IPS system or appliance to include:

1. A comprehensive, reliable flow of application, context and content information into the IPS system.
2. Advanced management and reporting capabilities to correlate and analyze data so organizations can respond quickly with the right measures.

This analysis suggests a three-layer model like the one shown in Figure 3. In this model, the Next Generation IPS is the core. But to provide all of the next-generation capabilities, it must be supported by two other independent but connected layers.

Figure 3: The Next Generation IPS and two supporting layers

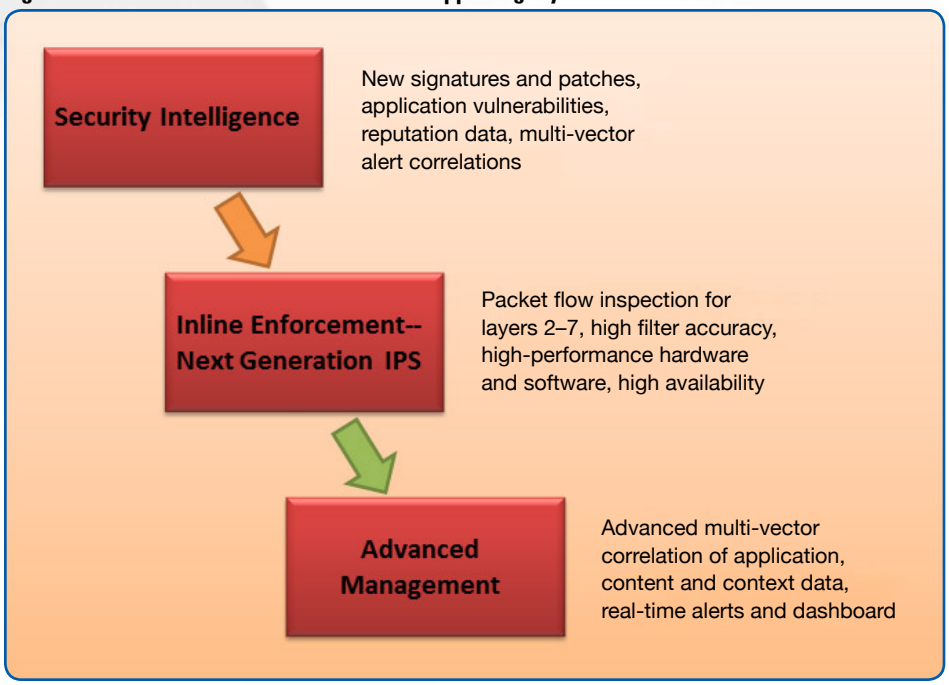

The upper layer is not a physical product at all, but rather a security intelligence service. This service supplies not only conventional threat signatures and patches, but also a rapid flow of context- and content-based information that can be used to detect and block new threats and threat types.

The middle, or inline enforcement, layer is the Next Generation IPS. This is an IPS with very high performance, application awareness, and the ability to detect and block network traffic based on context- and content-based information.

The bottom layer is an advanced management component that takes information from the IPS and system logs and correlates events related to applications, content and context. It then alerts administrators when specific combinations of events occur or when the quantity of suspicious events cross designated thresholds.

How a Next Generation IPS Might Have Stopped the RSA Breach

Can a Next Generation IPS stop advanced targeted threats?

There is no guarantee that a persistent attacker can always be thwarted. But we *can* show how a Next Generation IPS might have prevented the RSA breach. In fact, more contextual information and content awareness and other characteristics of a properly configured Next Generation IPS could have detected or blocked suspicious activities at seven different stages of the attack. This is illustrated in Table 1.

Table 1: Attack steps, Next Generation IPS countermeasures and results

Attack Step	Countermeasure from the Next Generation IPS	How the Countermeasure Might Have Prevented the RSA Breach
Spearphishing (1)	**Context Awareness:** Block inbound traffic from "known bad" web sites, including known spam and phishing sites and botnet controllers. **Content Awareness:** Detect spearphishing techniques (e.g., spoofing the sending address or links where the displayed text does not match the actual URL), then block spearphishing emails.	The spearphishing emails would not have reached the employees.
Malicious email attachment (2)	**Content Awareness:** Use content filters to identify and block emails with malicious attachments.	The emails with malicious attachments would not have reached the employees.
Exploit an application vulnerability to plant a Trojan (3)	**Vulnerability Protection:** Use filters to detect code designed to exploit a known vulnerability.	The emails with malicious attachments would not have reached the employees.
Trojan establishes a command-and-control link to an external system (4)	**Context Awareness:** Block outbound traffic to "known bad" web sites. **Content Awareness:** Detect outbound traffic characteristic of hacker command-and-control activities, then block that traffic.	The attacker would not have been able to contact the Trojan.
Attacker probes and scans the network (5)	**Context Awareness:** Detect network scans and share enumeration, quarantine the system initiating the scans, and notify system administrators and affected end users. **Context Awareness:** Use geo-location information to detect when the source of attacks shifts from external to internal, then quarantine the internal system.	The attacker's reconnaissance activity would have been detected and stopped.
Extract data from the target systems and move it to an internal staging server (8)	**Event Correlation:** Correlate alerts and events to identify deviations from normal activity (e.g., data transfers during non-business hours, data moved from departments with confidential data to departments that don't use that data and employees accessing systems they don't normally use), then alert system administrators.	The attacker's activities would have been detected before data was staged for exfiltration.
Transfer encrypted files to an external server (9)	**Context Awareness:** Detect and block outbound traffic to "known bad" domains and unapproved geographies.	The attacker would not have been able to exfiltrate the confidential data.

Protecting Against APTs: The HP Enterprise Security Network Defense System

The Next Generation IPS described above is not just a vision. On the contrary, there is one that has been tested and proven in some of the most demanding IT environments in the world. It is the core of the HP Enterprise Security Network Defense System (NDS), which is shown in Figure 4.

Figure 4: The HP Enterprise Security Network Defense System

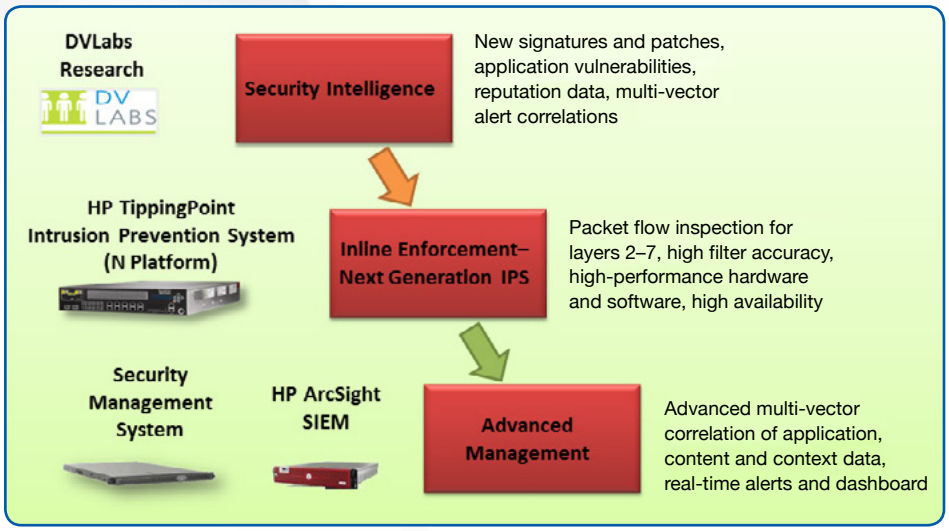

Security Intelligence

The security intelligence layer of the NDS is provided by DVLabs, HP's research organization for vulnerability discovery and analysis.

At DVLabs, teams of information security experts:

- Discover and verify zero-day attacks.
- Create a comprehensive set of threat signatures, patches and filters by monitoring and analyzing global Internet attacks and reverse-engineering malware.
- Work with customers and security organizations like SANS, CERT and NIST to compile one of the industry's largest and most reliable database of "known bad" and suspicious IP addresses and DNS names.
- Develop application filters to detect and block specific applications like social networking, IM, P2P, streaming video, online gaming and tunneling, and to limit functions embedded within applications, such as chat or file upload within Facebook and YouTube.

DVLabs packages this information into Digital Vaccines that continuously supply intrusion prevention systems with signatures, patches, context information and content-based rules that a Next Generation IPS can use to detect and block new threats.

In the context of the RSA breach, Digital Vaccines from DVLabs would probably have included information on:

- The (bad) reputation of the source of the emails.
- Phishing techniques used in the spearphishing emails.
- The malicious attachment containing the Poison Ivy Trojan (DVLabs has a filter to detect the Poison Ivy Trojan and its command-and-control traffic).
- The Adobe Flash vulnerability used to plant the Trojan (DVLabs released a filter to block this vulnerability on March 16, 2011).
- The reputation of the external system used for command and control of the attack.
- The reputation of the external system that was the destination of the files exfiltrated via FTP (in fact, DVLabs identified mincesure.com as a "known bad" site on March 17, 2011).

Inline Enforcement — The Next Generation IPS

At the core of the HP Network Defense System are the HP TippingPoint next-generation IPS appliances. These go beyond the capabilities of first-generation IPS to offer dramatically higher performance, high availability, and the ability to scan and interpret layers 2–7 (including applications). They provide the ability to detect attacks on many types of targets, including network devices, operating systems, enterprise applications, web applications and virtualization software.

These systems can also use information supplied by the DVLabs Digital Vaccines to enforce rules based on context, content and applications. In the context of the RSA breach, the TippingPoint systems might have been able to:

- Identify the emails as coming from a "known bad" web site and block them.
- Identify phishing techniques and block the emails.
- Identify the attachment as containing a Trojan and/or identify the attachment as including code designed to exploit the Adobe Flash vulnerability.
- Detect traffic to and from the external command-and-control system and block it.
- Detect the network scan, share enumeration and other internal reconnaissance activities of the attacker and alert system administrators.
- Detect the shift in attack patterns from external to internal and provide that information to the SIEM system for correlation and analysis.
- Detect traffic to the "known bad" mincesure.com site and block that traffic.

Advanced Management

HP TippingPoint Security Management System (SMS) appliances provide management capabilities to support Next Generation IPS environments. They give security administrators a continual view of security event logs to facilitate immediate cyberattack containment, perpetrator location and identification, and damage mitigation. They also compile log information to present big-picture analysis and trending reports, real-time graphs on traffic statistics and filtered attack events reports.

These management features can be supplemented by the HP ArcSight SIEM (Security and Event Management) appliances, which provide extremely sophisticated capabilities for correlating and analyzing many types of security events.

In the context of the RSA breach, the SMS and ArcSight SIEM systems could have been used to detect deviations from normal activities, including the attacker accessing the target systems from compromised PCs and data transfers from the target systems to the internal staging server.

Conclusion

The March 2011 attack on RSA showed how vulnerable even security-conscious companies are to today's well-funded and patient cybercriminals and malicious state actors. It also illustrated why the ideas in Pescatore and Young's "Defining Next-Generation Network Intrusion Prevention" research note are so important.

Technology alone will never guarantee success against cybercriminals. But a close look at the steps of the RSA breach and potential countermeasures shows that a Next Generation IPS with application awareness, content awareness and context awareness gives organizations powerful new ways to protect against sophisticated threats.

For more information about HP Enterprise Security and HP TippingPoint Next Generation IPS solutions, please visit www.hpenterprisesecurity.com/ngips.

[i] Wikipedia entry for advanced persistent threat: http://en.wikipedia.org/wiki/Advanced_persistent_threat; *Under Cyberthreat: Defense Contractors*, Bloomberg Businessweek, July 6, 2009: http://www.businessweek.com/technology/content/jul2009/tc2009076_873512.htm; *Advanced Persistent Threat, What APT Means To Your Enterprise*, Presentation by Gary Hoglund: http://www.issa-sac.org/info_resources/ISSA_20100219_HBGary_Advanced_Persistent_Threat.pdf; also see http://www.businesswire.com/news/home/20120119005051/en/HBGary-HP-Enterprise-Security-Partner-Deliver-Advanced.

[ii] *Open Letter to RSA Customers*, Art Coviello, executive chairman of RSA: http://www.rsa.com/node.aspx?id=3872; *RSA SecurID Breach Cost $66 Million*, InformationWeek, July 28, 2011: http://www.informationweek.com/news/security/attacks/231002833.

[iii] The discussion of the RSA breach is based primarily on information from these sources: *Anatomy of an Attack*, RSA "Speaking of Security" blog: http://blogs.rsa.com/rivner/anatomy-of-an-attack/; *How We Found the File That Was Used to Hack RSA*, F-Secure "News from the Lab" blog: http://www.f-secure.com/weblog/archives/00002226.html; *SC Congress Canada: RSA security architect discusses SecurID breach*, SC Magazine: http://www.scmagazine.com/sc-congress-canada-rsa-security-architect-discusses-securid-breach/article/205402/; *Researchers Uncover RSA Phishing Attack, Hiding in Plain Sight*, Wired Threat Level web site: http://www.wired.com/threatlevel/2011/08/how-rsa-got-hacked/.

[iv] *Defining Next-Generation Network Intrusion Prevention*, John Pescatore and Greg Young, Gartner, October 7, 2011, ID No. G00218641.

