IBM X-Force Threat Intelligence Quarterly, 4Q 2015

Explore how cybercrime trends and incident forensics are evolving—based on the real-world insights of the IBM Security Services team

IBM

Contents

2 Executive overview

3 The top four cybercrime trends

13 The power of "indicators of compromise" for incident forensics

18 About X-Force

19 Contributors

19 For more information

Executive overview

In this final issue of 2015, IBM® X-Force® shifts the focus to our in-house experts at IBM Security Services. The IBM Security Services team has an extensive global reach and experience with addressing cyber-security concerns and incidents affecting clients across a broad range of industries. Insights garnered from each client engagement are collected, forming a vast pool of knowledge that our analysts draw upon to identify common threads. These threads are woven together to form a picture of current security trends, techniques and tools used by cybercriminals, and reoccurring gaps in our clients' security postures.

Many of the security incidents to which the IBM Emergency Response Services (ERS) team responds involve fundamental breakdowns in sound security practices—that is, they could be prevented. While the incidents occur around the globe, many of them share certain characteristics and fit recurring patterns. Our report begins by highlighting four key trends the ERS team has observed throughout 2015, including an increase in "onion-layered" security incidents, ransomware attacks and insider threats, and the transformation of security issues into a boardroom priority.

These security trends yield an important question: How can your enterprise find the footprints that attackers leave behind when they breach your defenses? Fortunately, the attackers aren't the only ones who have a collection of tools at their disposal. Our second article addresses indicators of compromise (IOCs), which provide the digital evidence that an attack may have occurred and are an important tool in forensic analysis following a breach. We focus on how security teams can use IOCs to track advanced attackers, assess the level of compromise and remediate issues before significant damage occurs.

It's been a tough year for security teams. Insider threats, malware, stealthy tools and morphing attacks continue to challenge organizations of many sizes in 2015. When IBM X-Force looks back across the year, we see many areas for improvement. The good news is that organizations can take stronger responsibility, make a few small changes and see a big impact for the long term.

The top four cybercrime trends

What kinds of security incidents are striking again and again, across every industry? Get an insider's look from the IBM Emergency Response Services team.

Security incidents have been on the rise for the past few years, and most experts in cyber security believe the trend will only continue to intensify. Here, though, our subject is not the high-profile, headline-grabbing attacks we all know about but the everyday struggle of organizations everywhere, in every industry, to protect their data in a world of thieves.

This article must begin with a basic fact of cyber-security life. Many, if not most of the security incidents to which the IBM ERS teams respond involve fundamental breakdowns in sound security practices, and are wholly preventable.

Our goal here is to provide insight into the security issues we've observed. While the incidents that we respond to around the globe are diverse, many of them share certain characteristics and fit recurring patterns. This report highlights four key trends we're watching in 2015:

- "Onion-layered" security incidents
- Ransomware attacks
- Attacks from inside an organization
- Greater management awareness of the need to address security threats proactively

About this article:

This article was created by IBM Emergency Response Services, a team of highly skilled security consultants who work hand in hand with clients to help them prepare for, protect against and respond to security incidents. Based on the collective experience of these consultants working in the field, this report is intended to provide you with deeper insight into current threats and security events.

Trend 1: Onion-layered security incidents on the rise

As the name suggests, an "onion-layered" security incident is one in which a second, often significantly more damaging attack is uncovered during the investigation of another more visible event. The security team has to carefully peel back layers of forensic information in order to determine the root cause of each event under scrutiny. Usually, the actors involved in the two incidents are:

- The *script kiddie*, an unsophisticated attacker who launches highly visible attacks and is careless about getting caught. The script kiddie is typically responsible for the original incident that alerts security teams to a security problem, which can lead to the involvement of the ERS team.
- The *stealthy attacker*, a much more sophisticated and careful hacker who might go undetected for weeks or even months while expanding his or her grip on the victim's network. The second, more serious incident is this actor's work.

Of all the incidents that the ERS teams encountered, these complex, multi-layered attacks were the most demanding of investigative time and resources to ascertain the facts, find the root causes, develop a timeline of events, and provide the client with recommendations on how to resolve the issues that allowed the attackers to get into their network.

Discovery of the incidents could begin with someone calling an organization's support number to report that a website had been defaced, or with system administrators noticing high CPU usage or an unusual amount of traffic coming out of a server, or discovering unusual files on that server. The subsequent investigation would show that this was the work of an attacker (the script kiddie) who had compromised the server by exploiting some long-standing vulnerability or configuration error. The common trait among a number of compromised systems we investigated was that they were running old operating system versions that hadn't been patched in a long time.

While analyzing the compromised system, however, some fact totally unrelated to the attack under investigation would emerge and immediately widen the scope of the incident. For example, the team might find that the source of a backdoor on the analyzed system wasn't the Internet but another of the client's servers. Or the secure socket shell (SSH) logs would reveal logins from suspect countries where nobody managing the server was located, and those logins had been happening for months before the initial security incident was raised. When the scope of the investigation was widened and more systems included, a whole new story would emerge: a second group of attackers, far more sophisticated and stealthy than those initially identified, had been compromising servers for months and, in some cases, had managed to jump from Internet-facing servers to the internal network.

The tools and techniques used by this second group differ substantially from those used by the script kiddies. Their goals differ too. The script kiddies scour the Internet for "low-hanging fruit," the servers that can be compromised quickly and easily, and they use them for a limited time to send spam and scan other servers on the Internet. Or they deface the website and move on to other targets once they are discovered. These script kiddies give little thought to covering their tracks.

In contrast, stealthy attackers might gain access to a system by exploiting the same vulnerability as the script kiddies, but they use a far more sophisticated combination of commercial tools, malware/rootkits and backdoors to increase their access level on the client's network and compromise additional systems over several weeks of expansion. They cover their tracks much more effectively and use anti-forensic techniques such as manipulation of timestamps to remain undetected for as long as possible. What's more, their ultimate goal can be far more unsettling—ranging from data theft to corporate espionage or worse.

Although these sophisticated attackers go to great lengths to remain undetected—and are successful in many cases—the team identified examples of telltale signs that can lead to early discovery of malicious activity, including:

- **Alerts generated by anti-virus software about Trojans or hacking tools on Internet-facing web servers.** Anti-virus tools can be set to automatically delete Trojans, leading to a false sense of safety. The important question—how did the Trojan get on a server in the first place—may remain unaddressed. Knowing the "how" helps determine the appropriate mitigation to implement, in order to prevent future infections. In fact, anti-virus alerts can indicate a stealthy attacker's first attempts to compromise a server.
- **Servers rebooting unexpectedly or other unusual behavior.** When a server acts up, troubleshooting typically focuses on fixing the problem. The security-aware team will go the extra step to ask how the problem started in the first place. For example, an IT team finds some unknown software on a server that is causing unusual behavior. Removing it fixes the problem. Just as important is determining how the software got there in the first place. It could be a malware kit that is affecting the server in question and others in the network—a sure sign that a security breach has likely taken place.
- **Suspicious log records.** It is important to examine the root logins recorded in SSH logs. Two signs of trouble can crop up here. For one thing, authorized users should not be logging in as "root" directly from the Internet. This can indicate that SSH settings have been tampered with. The location of the IP address can be another sign of unauthorized access. Be suspicious if a country of origin is not where legitimate administrators are located.
- **User lockouts.** Numerous users getting locked out and asking for their accounts to be reset should be a red flag for administrators to alert their security teams. This can be a sign that someone has used techniques such as dictionary or brute force attacks to "crack" user passwords. Left undeterred, the attacker's next steps could be to scan Internet-facing servers to find other locations where they can use the stolen credentials.

What factors can facilitate these types of attacks?

Incidents like those described above can result from two main issues:

- **Breakdown 1: Old and unpatched systems exposed to the Internet.** In many of the cases we worked on, the initial break-in was to an Internet-facing server running an old, often out-of-support operating system that had not been patched in years. This can indicate incomplete patch management procedures and a general lack of oversight of the systems deployed in the network.
- **Breakdown 2: Clients having little visibility into their network.** Clients victimized by onion-layered attacks were often not watching what was happening in their network. Typically they ran anti-virus software, used firewalls and sometimes employed an intrusion prevention system—but little was being done to manage alerts generated by these security devices. In cases like this, only major service interruption issues were addressed and investigated to any extent. This lack of visibility can make an organization an easy target where malicious activities could go undetected for a long time. Were it not for the disruptive event caused by the script kiddies, the client might never have noticed anything wrong.

What is the impact?

While the recovery of systems compromised by script kiddie attacks might take only a few days of an operation team's time and effort, the job of finding a root cause, then fully understanding and remediating the work of the stealthy attackers could take months. Meanwhile, the stealthy attacker could roam the network undetected, ultimately trying to gain access to the client's "crown jewels."

How can your organization be prepared?

- **Keep systems updated.** Take the time to test updates and apply them quickly. This includes keeping the operating system version up to date. If you have old systems that aren't updated regularly, don't expose them to the Internet.
- **Increase your visibility into what's happening on the network.** This can be done by adopting a combination of products that provide intrusion protection, security information and event management (SIEM) and network traffic monitoring capabilities.
- **Build an internal security operations center** (or outsource it to a managed security services provider) to monitor the alerts and events generated by your security systems, and follow up and investigate all "odd findings."
- **Create operational procedures** for responding to common events such as server reboots, account lockouts and alerts generated by the anti-virus software. An event happening on an internal workstation may need to be handled differently than the same event on an Internet-facing server.
- **Make sure the level of logging is appropriate** and that logs are stored centrally to make them hard to tamper with and easy to access during a security incident.
- **Periodically perform penetration testing exercises** to identify systems and applications showing vulnerabilities that have to be addressed quickly.

Trend 2: The year of ransomware

The infection scenario most commonly encountered by ERS in 2015 was ransomware. As its name suggests, this is a kind of malware that steals something from the user and demands a ransom to give it back. Ransomware can be divided in two broad families:

- The first family simply locks the system and tricks the user into thinking that unlocking it requires paying a ransom. This is the less dangerous kind of ransomware, since no actual harm is done to the infected system and no information is lost.
- The second family actually encrypts files on the system's hard drive. Instructions on how to pay the ransom and get the key to decrypt the files are left in text files disseminated on the hard drive. This is the more dangerous kind of ransomware, since breaking encryption often isn't feasible and might result in losing information even if the ransom is paid. A particularly destructive variant of this second family will encrypt not only files on the hard drive of the infected computer, but also network shares, potentially targeting the files of the user's organization.

A widespread belief in the computer security industry is that ransomware is a profitable underground business, and most vendors predict that it will remain a common threat through the end of 2015 and beyond, migrating to mobile devices as well. Latest evolutions include malware that encrypts specific fields of a database associated with a web application.[1] This is done by introducing malicious code into the application—code that will encrypt and decrypt the data being inserted or being extracted on the fly, as requested by the application.

By letting the malicious code run for weeks or months, then removing the key to encrypt the data, the attacker ensures that parts of the database will be encrypted with no way to decrypt them. At this point the web application will stop working and the attacker will demand a ransom for the key to decrypt the data. This kind of ransomware attack, where even backups won't help in restoring the encrypted data, is quite dangerous. Whether or not it will become more widespread is unclear as yet.

What factors facilitate the attacks?

For ransomware to succeed, attackers rely on a multitude of security and procedural breakdowns. In some cases, clients had recurring infections during the year. This was because, although some of the factors leading to infection were addressed and resolved, nothing was done to resolve the fundamental breakdowns that facilitated the initial infection:

- **Breakdown 1: Not backing up data.** When a client has a ransomware issue, one of the first questions the ERS team asks is, "Do you have backups of the encrypted files?" All too often, the response is "no." If your organization is panicking because vital files may be lost, it's time to reevaluate your backup methodology.
- **Breakdown 2: Poor patching procedures.** Frequently, the ERS team is asked to find out how ransomware was able to enter a client's environment. And often the answer is inadequate patch management. High-severity software patches that should be applied within hours are sometimes applied months later, or not applied at all. A well-known infection vector of ransomware can exploit unpatched operating system vulnerabilities to give attackers access to the system resources they want to lock or the data they want to encrypt.
- **Breakdown 3: Lack of user awareness.** Many security professionals believe that users are the weakest link in the organization. If users are not aware of safe computing practices, they can inadvertently undermine significant investments in information security just by clicking on the wrong link or visiting an insecure website. ERS teams have repeatedly observed a lack of user awareness as a key shortcoming during ransomware-related engagements. A well-trained workforce is a very inexpensive multiplier for an organization's security investment.

What is the impact?

The impact varies depending on the organization's size and level of preparedness. Some lose key intellectual capital. Others are more fortunate, suffering only operational disruptions lasting days or weeks. In the worst cases, mostly among small to midsized businesses, ransomware attacks can be devastating, causing a complete shutdown of business.

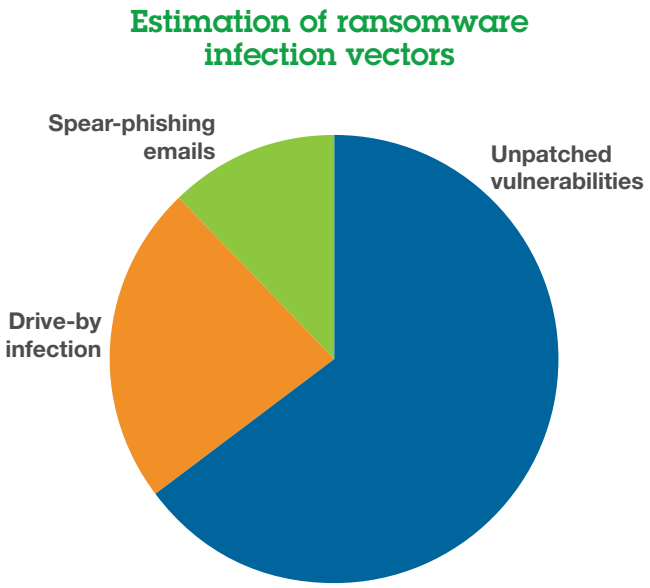

Figure 1. The IBM ERS team estimates that three primary vectors are the source of ransomware infections.

How can your organization be prepared?

Due to the attack vectors this threat exploits (see Figure 1), the most effective long-term strategy is to focus on improving both patching procedures and safe computing practices. Create a company-wide training program on safe computing practices. For example, every employee needs to know how to recognize the signs of phishing attempts.

Before clicking, everyone should know to ask themselves:

- Is the source of this email or communication reliable?
- Did I ask for this attachment or link?
- Does the link I received for company X take me to their normal website?

If the answer to any of these questions is "no," users must have a quick, easy way to report the email as suspicious. This allows the organization's security team to spot recurring trends and identify attempts at mass attacks.

Other preventive measures can supplement the security training:

- Anti-phishing techniques such as checking email headers on the mail server can help block phishing attempts and prevent phishing emails from reaching the intended recipients.
- Should all other countermeasures fail, software designed to catch anomalies related to binaries, processes and connections can also help identify many kinds of malware, ransomware included.

Security practices beyond user awareness will help in recovering from ransomware incidents and performing impact analysis:

- Configure your anti-virus software to quarantine malicious files instead of deleting or cleaning them. Then they can be analyzed later if needed.
- Review the need for open sharing between networked endpoints, and disable as many as possible to limit the attack surface available to ransomware.
- Make sure that backups are created and tested regularly. This will go a long way to help quick recovery from ransomware incidents, minimizing information loss and recovery time.

Should backups not be available in the recovery phase, there might be other ways to recover data:

- File recovery software or professional services can be effective with ransomware variants that make a copy of the files before encrypting them and then deleting the original. Success hinges on the frequency with which content changes on the hard drive and ultimately on how much time elapses between detection of the problem and the attempt to recover files.
- Microsoft Windows Volume Shadow Copy Service can help. Volume shadow copies are usually deleted by the malware upon encryption of the files in an attempt to thwart recovery attempts, but sometimes the deletion fails and recovery of the malware-encrypted files is a possibility.

In the containment phase, enabling logging on critical folders and files will help quickly determine the initial point from which the ransomware spread and suggest ways to contain the ransomware at the network level. Thorough logs and network data are also critical for determining the extent of the damage. Logging on critical files and folders should monitor:

- Which user accesses them
- Which user changes them
- From where in the network they are accessed

Trend 3: Malicious insiders on the attack

During 2015, the ERS team was called on several times to assist with unexplained network outages—both to stop the outage and find the root cause. The symptoms ranged from routers that had their configurations erased to firewalls with unauthorized rule changes. In some cases the impact was only temporary and resolved within a few hours without intervention, but the problem would keep reoccurring over time.

Due to the sometimes volatile nature of these issues and the difficulty of distinguishing their true nature from "normal" service outages, some of the situations went on for weeks before it became clear that a security incident needed to be declared and the ERS team was engaged.

In the best-case scenario, investigations showed that the changes were due to the use of a shared administrative account, but their real source was difficult to determine. In the worst-case scenario, no logs were available, so finding the cause of the outage was impossible.

A series of patterns emerged from the ERS team's investigations:

- There were shared accounts with administrative privileges.
- Password sharing between team members was not discouraged.
- Passwords were routinely set to never expire.
- Passwords were "easy."

The common thread is that accountability was not enforced. Bad password policies seriously compromised the efficacy of termination procedures. Whenever a system or network administrator left the organization, disabling their personal accounts did not limit their ability to perform unauthorized activity on the network via one or more of the shared accounts they had routinely used in their job. As a result, ex-employees with ill will toward former employers held powerful weapons they could use to express their resentment. They simply needed a way to get back into the network.

In most malicious insider attacks we've seen, the disgruntled employee typically "prepared for departure" by installing remote administration tools (RATs) such as LogMeIn or TeamViewer for access to the employer's network. Such tools only establish outbound connections to the Internet, so they are rarely monitored or blocked by a firewall. In many cases the RATs were installed on several servers. Sometimes a valid (and shared) virtual private network (VPN) account was also used, and the employee would change the means of connecting to the network when one of the shared accounts was discovered.

With these pieces in place—one or more shared accounts, an administrator's knowledge of the network, and a way back in such as a valid VPN account or a RAT—an embittered ex-employee could cause a lot of damage for a long time.

What factors facilitated the attacks?

- **Breakdown: Lack of accountability.** Shared accounts and a lack of accountability were the main issues. Routinely implementing and using shared accounts made termination procedures highly ineffective.

What was the impact?

The actions of a malicious insider can cause disruption of normal operations and potentially other harm. Even if the damage isn't persistent, countless hours of troubleshooting can be spent by an organization's IT operations team to investigate and fix the issues caused by the disgruntled employee.

How can your organization be prepared?

Knowledge can't be stripped from an employee leaving an organization, but there are ways to minimize the risk of that knowledge being used for malicious purposes:

- Enforce accountability and good password policies.
  - All administrators should have their own username and password and always use them to perform normal administrative tasks. This rule should apply to all employees, but it's critical for those with administrative permissions on the network or infrastructure.
  - Password sharing between team members should be prohibited.
  - If prohibiting shared administrative accounts is not an option, they should be limited to the bare minimum. Their usage and the activity performed by them should be monitored closely.
  - Passwords should be reset periodically.
- Termination procedures should be enforced.
  - All credentials for an employee leaving the organization, voluntarily or otherwise, must be disabled immediately upon termination.

That is the main set of policies to be enforced at all times. Other recommendations may help in forensic discovery and analysis of an incident:

- All network devices and servers should have their times synchronized with a common Network Time Protocol (NTP) server. This is to ensure that timestamps of the logs are consistent and can be correlated.
- An appropriate level of logging should be enabled on all servers and network devices. Information recorded by logs should include at least:
  - Time of login
  - Account used to login
  - Source of login
  - Switches between users (for example, user X switching to super-user root)
  - Activity performed by user (preferable)
  - New account creation, particularly super-user accounts
- To avoid the possibility of an attacker tampering with such logs, they should be stored centrally on a server dedicated to their preservation. A syslog server would provide the bare minimum required. A SIEM system that provides added features such as correlation between events and generally enhances oversight of the network is preferable.

Remediation in the containment phase of an incident

Commercial RATs such as LogMeIn and TeamViewer are the malicious former administrator's usual means of guaranteeing access to the network. Most of these tools work very like TeamViewer:

- There is a RAT client installed on the attacker's computer.
- There is a RAT server installed on the target computer where the attacker wants to connect. This will be one of the servers in the client's network.
- There is a rendezvous or master server on the Internet, managed by the company developing the RAT software.
- Both RAT client and server establish TCP connections to the master server on the Internet. This creates a virtual connection between client and server that allows the attacker to remotely control the computer running the RAT server, as illustrated in Figure 2.
- Client and server can adapt to changing network environment and firewall rules in place when connecting to the master server, making it very difficult to block such connections.

Should unauthorized RATs be detected during an investigation, or even deemed likely, a very effective remediation is to use the domain name server to block access for the master servers of all known RATs, such as teamviewer.com, master*.teamviewer.com, logmein.com or gotomypc.com. Unfortunately, new RATs appear frequently and existing ones change their infrastructure and add or remove master servers on a continual basis, limiting the effectiveness of this measure.

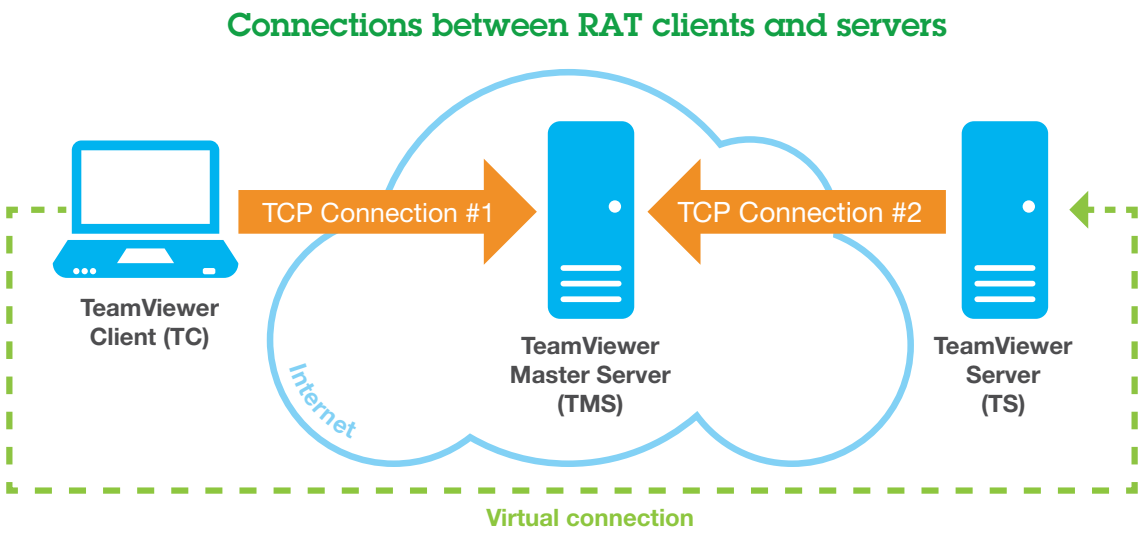

Figure 2: Internet connections established between the RAT client—TeamViewer, in this illustration—and the RAT server allow an attacker to remotely control the computer running the RAT server software.

Trend 4: Greater management awareness of security problems

In recent months, the ERS team has observed that people in positions of oversight—management, boards of directors, audit committees—are asking more questions about their organizations' security posture. Given the recent high-profile breaches of many well-established organizations, this is a welcomed trend. ERS clients today are asking about:

Mock tabletop exercises

Tabletop exercises are a great way for organizations to prepare for a security emergency. The ERS team has facilitated a wide range of mock tabletop exercises for clients, including stress tests, educational scenarios, technical and non-technical discussions, and cross-functional reviews. For many clients, this is their organization's first attempt at conducting any sort of mock exercise.

Incident response plans

Organizations are placing greater emphasis on planning for computer security incidents. Many recognize that security threats, despite considerable investments in protection and prevention, are inevitable, so creating the ability to respond quickly and efficiently may mean the difference between a short-duration event with limited impact and a long-running disaster. Driven largely by management interest, organizations have been creating incident response plans. Those with plans already in place have been asking ERS for third-party reviews to bolster strongpoints and identify weaknesses.

Enterprise information system risk assessment

Aware of the high potential costs of a security breach (see Figure 3), management is pushing ever harder to get ahead of the threat curve. Many information security techniques focus on detecting malicious software or actors already within an environment. ERS clients eager to reduce their overall risk footprint are now asking for environmental assessments to look for risk factors a malicious actor might exploit—information systems running unknown processes or communicating to foreign systems, for instance—so the risk factors can be mitigated.

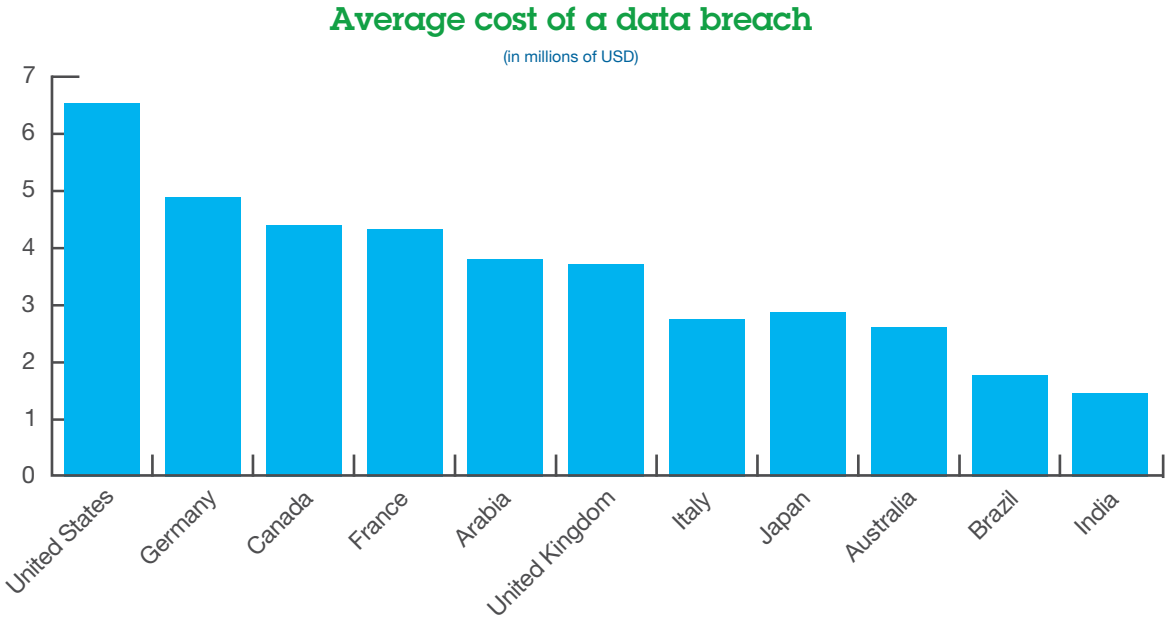

Figure 3: An individual data breach can cost on average up to USD6.53 million, according to the 2015 Annual Cost of Data Breach Study: Global Analysis from Ponemon Institute (sponsored by IBM).

Conclusion

Organizations today are going back to the basics. The major cyber-security trends of 2015—the challenge of recognizing stealth attackers on the network, ransomware, malicious insider attacks and growing management attention to enterprise security readiness—can largely be addressed by focusing on "security 101." Think patch management, user education, proper password procedures and standard security practices. A defense-in-depth strategy built on these components will help organizations reduce the risks we see today and expect tomorrow.

Readiness is the key. Reduce enterprise risk to limit attackers' opportunities as much as possible, but understand that attacks will still come and organize your defenses to react quickly and cohesively when they do. Expert help is useful in that endeavor. Experienced professional incident response consultants can suggest the most effective way to contain, eradicate and recover from an attack, pinpoint the root cause and take action to help prevent it from happening again.

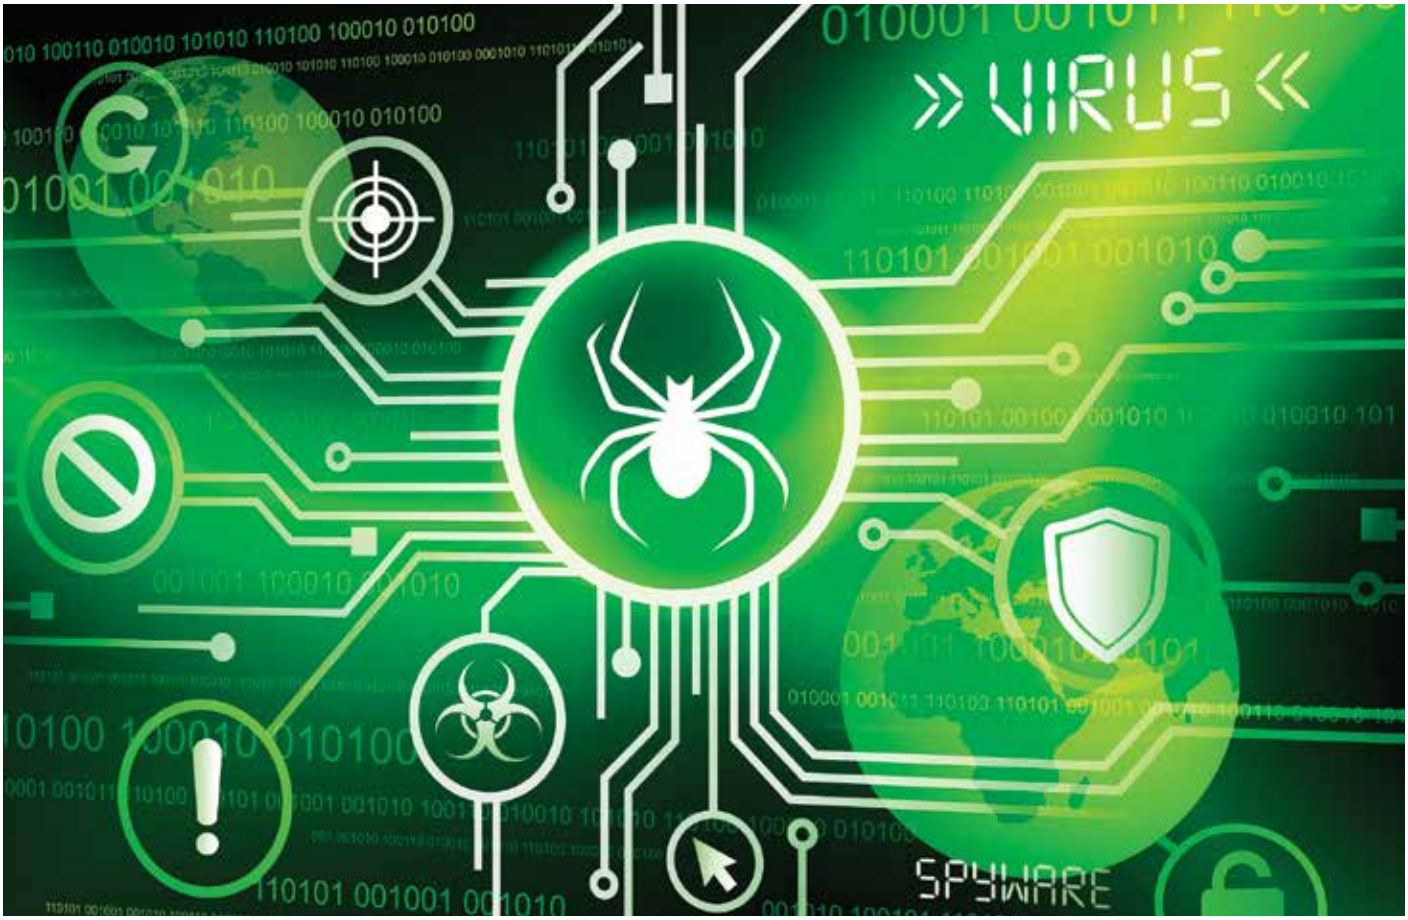

The power of "indicators of compromise" for incident forensics

Attackers often leave digital evidence all over your network. Learn how to detect the intrusions and be prepared to respond.

Indicators are everywhere. The "check engine" light tells you when one of your car's systems has failed. Your cell phone alerts you when the battery is low. Your home security system sounds an alarm if it detects an intruder, and your home computer displays a warning message when a device or piece of software malfunctions. From a design perspective it seems simple: you understand what to look for and you design a monitoring control around it. But what if your task is to reliably detect intrusions within a network or operating system? What if you're building a system to identify with high confidence artifacts that indicate an intrusion? That's not simple at all.

The term "indicator of compromise" (IOC) was first used by government organizations and defense contractors attempting to identify advanced persistent threats (APTs). Since 2007 the term has been commonly used throughout the information security industry. IOCs are digital evidence that suggest an attack may have occurred and are an important tool in forensic analysis following a breach.

An evolution of the IOC is the "indicator of attack," or IOA. IOAs play a major role in identifying the intent of an attacker regardless of the malware or exploit used, and as a result, next-generation security solutions are moving to an IOA-based approach pioneered by CrowdStrike, an IBM partner. IOAs will be covered in a future IBM X-Force research paper.

Here our focus is on IOCs. Our goal is to illustrate their importance and help you better protect your enterprise network environment from advanced threats.

Indicators of compromise

Indicators of compromise are evidence on a computer indicating that the security of the network has been breached. Investigators usually gather this data after being informed of a suspicious incident discovered during a routine, scheduled scan or after the discovery of suspicious activity on the network. They gather the information to create smarter tools for detecting and quarantining suspicious files or blocking suspicious traffic.

Let's look at what IOCs are, exactly, and how you can leverage them to detect anomalies within your network.[2]

Unusual outbound network traffic

Patterns of unusual traffic leaving your network perimeter should always be investigated. Modern attack methods make keeping attackers out of a network difficult, but outbound patterns are much more easily detected. Command and control (C&C) traffic from compromised servers may be visible, allowing victims to respond before data is lost or damage caused.

Anomalies in privileged user account activity

Attackers often try to escalate privileges of a user account they've hacked. Monitoring privileged accounts for unusual activity not only opens a window on possible insider attacks, but can also reveal accounts that have been taken over by unauthorized sources. Keeping an eye on systems accessed, type and volume of data accessed, and the time of the activity can give early warning of a possible breach.

Geographical irregularities

Irregularities in login patterns can provide reasonable evidence of compromise. Connections to places where your organization does not normally do business might mean your sensitive data is being stolen. Accounts noted as logging in from multiple IPs in a short period of time paired with location tagging can provide enough evidence to take a deeper look at that activity.

Other login red flags

Excessive failed logins or attempts on accounts that don't exist are signs that an attacker is trying to guess credentials. Look specifically for login attempts with usernames of employees who wouldn't normally be working after hours. That might be a perpetrator at work, not the employee; it's a red flag for investigation.

Surges in database read volume

If an attacker penetrates your database storage, the exfiltration of that data, especially credit card tables, will generate a read volume well above normal for those tables.

Large HTML response sizes

An attacker using a SQL injection attack against your database will cause a larger than normal volume of HTML responses. For example, a 20 MB response to a query that is normally around 200 KB can indicate that the attacker has successfully executed a SQL injection attack and dumped the entire credit card or user account table.

Large numbers of requests for the same file

When an attacker finds a worthwhile target on your network, for example a vulnerable web application written in PHP, they will try multiple attack strings focused on a specific file. If you detect a single source creating a high volume of requests to a specific file, such as "anyfilename.php," you should be immediately suspicious.

Mismatched port-application traffic

Communications on non-standard ports could be an indication of foul play such as C&C traffic masquerading as "normal" application behavior.

Suspicious registry changes

Malware often persists across system reboots by modifying the registry to launch a startup process or to store operational data. Always create a clean baseline registry snapshot and monitor for changes to this "template" that could indicate a registry-based IOC.

DNS request anomalies

A large spike in Domain Name Service (DNS) requests from a specific host can indicate possible malicious activity. Watch for patterns of DNS requests to external hosts, and compare them against geographic region and host reputation data. Filtering solutions that are tied to threat intelligence tools can help detect and mitigate malware by discovering that it is communicating with its C&C infrastructure.

Unexpected patching of systems

Patching systems is one of the most normal transactions that can occur on a network, but the patching of critical systems out of cycle could indicate malicious activity. When attackers compromise a system, they want to make sure no other group compromises it, so they patch and harden it to prevent other attackers' access.

Bundles of data in the wrong places

In many cases attackers store large amounts of compromised data prior to exfiltrating it. They try to hide it in unusual places, such as the root directory of the recycle bin on a Windows-based server or directories on Linux machines that contain temporary files or cached data.

Web traffic with superhuman behavior

Infected machines compromised by click-fraud campaigns can generate high volumes of web traffic far faster than users sitting at a browser possibly could. In corporate networks where the users are required to use a prescribed browser, watching for user agent strings that don't match that internal mandate can help identify malicious web traffic.

Searching for indicators of compromise

Hunting for IOCs makes it easier to track advanced attackers.[3] The defense-in-depth lifecycle (see Figure 4) provides a roadmap. Consider performing some or all of these steps.

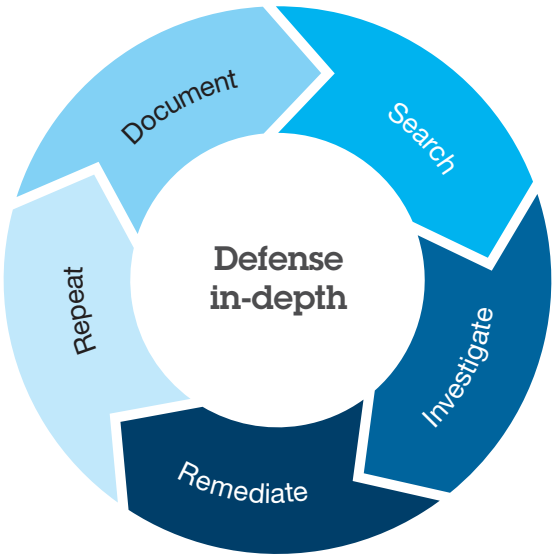

Figure 4. The "defense in-depth" lifecycle approach to tracking advanced attackers.

Step 1: Document attack tools and methods

- Profile your network traffic patterns to get a sense of what's normal. Focus on main protocols, especially the ones used by attackers, such as DNS and HTTPs.
- Collect and examine log file entries. Tools like log management and SIEM systems can automate much of this effort and provide an interface for visualizing data patterns and detecting suspicious activity.
- Leverage metadata to hunt for IOCs.
- Subscribe to IOC data feeds from organizations that analyze malicious tools and keep an up-to-date repository.

Step 2: Use the harvested intelligence to search for attacker activity

Configure your security defense tools to look for attacker activity using the data you gathered in step 1, including IOCs and deviations from normal behavior. These configurations may include blocking or alerting on:

- Activity from suspect IP address ranges or geographies with a poor reputation for hosting attacks (IP reputation).
- Attempts to exploit vulnerabilities: often intrusion prevention systems (IPS) and endpoint security systems will issue alerts on patterns that indicate exploit activity and may even identify specific vulnerabilities and known exploits.
- Hashes of known tools in the attacker arsenal: one of the first things attackers will do after gaining a foothold in a victim's environment is upload their toolkit so they can continue to infiltrate the organization.
- New usernames created locally.
- Usernames that were probed on other systems.

Step 3: Investigate security incidents and assess the level of compromise

- Begin with what is obvious: system IP, DNS, user, timestamp. Determine the number of systems or applications that are affected, the number of attempts to access the system or the application, and the degree of penetration the attacker achieves.
- Establish a timeline to determine if other events occurred. Examine all files with time stamps (logs, files, registry); the content of email communications and messages; information about system logon and logoff events; indications of access to specific Internet documents or sites; and the contents of communication with known individuals in chat rooms or other collaborative tools. Note that some of these may be subject to corporate policies or local laws (or both) regarding privacy protection. Check your individual organization's policies before proceeding.
- Check for evidence of document destruction.
- Search for incident-specific IOCs such as exhibiting patterns within working directories or using particular hosts and accounts. Use available tools like IOC Finder[4] to assist with your searches.

Step 4: Remediate

Identify:

- Compromised hosts and user accounts.
- Active (beaconing) and passive (listening) points of exfiltration.
- All other access points such as web servers, VPNs and terminal services.

Perform the following:

- Reset passwords.
- Remove points of exfiltration.
- Patch vulnerable systems being exploited for access.
- Activate your incident response team.
- Continue searching for IOCs to ensure remediation tactics are successful.
- Set trigger points to alarm if the attacker returns.

The art of IOC authoring

Practice creating IOCs with creativity in mind. The best IOCs have the following properties[5]:

- The IOC only identifies specific attacker activity that has been harvested due to its suspicious nature. For example, look for a specific file by MD5 sum (hash), file name, size, create date, or other file attributes. Look for a specific entity in memory (process information, running service information). Look for a specific entry or set of entries in the Windows Registry. Using these approaches in various combinations provides better matching ability and fewer false positives than searches for individual artifacts.
- The IOC is simple and evaluates information that is easy to collect and calculate.
- The IOC is difficult for the attacker to evade without changing tactics, tools or approach methodology.

Creating effective indicators

Unlike some other data standards used to describe threat information, there is no one-to-one mapping of an instance of a threat (such as a piece of malware) and the particular data standard used to describe it. An IOC of "OR filename = *.bat" is definitely going to identify a lot of files, but it's a rather poor indicator that will generate many false positives since it matches every executable batch file on a system. Better IOCs achieve the best true positive rate while having the lowest false positive rate (flagging things which are normally found on a system or not related to an intrusion).

More complex use cases and techniques combine variations of the following approaches:

- Instead of just looking for specific file artifacts in one part of the operating system or network, groups of artifacts can be combined using the logic of OpenIOC, an open framework for sharing threat intelligence, to create a match on artifact groups common across families of malware or other intrusion tools.
- Instead of hunting down a specific known bad file, an incident responder could make a whitelist of the files known to be in a directory, and then catch all the files not on that list. This is especially effective, and important, on critical systems with limited activities such as point-of-sales terminals, industrial control systems, and systems that contain credit card data, personally identifiable information (PII), or other sensitive data.
- Look for specific locations in the file system, registry, or other parts of the operating system that hostile actors regularly use in the course of their intrusion, even if this has nothing to do with the initial exploit or compromise.
- Look for sets of artifacts left by tools or toolkits used by adversaries that would be expensive for them to change or modify.
- Look for signs of adversary activity on systems used for lateral movement that were not directly compromised but show signs of activity outside their normal usage patterns.

In real-world cases, IOCs can combine any and all of the above types of functionality, or you can use just use a single type. Investigators tailor IOCs to the needs of the investigation, and the flexibility of OpenIOC allows them to do that as the case evolves, without having to write a new indicator.

IOC sharing and detection tools

Rapid communication of threat data makes it possible to quickly identify IOCs and defend against attacks. Interest in collecting and storing IOC-based threat information is high, so your network security professionals can leverage several newly-developed free or open-source tools to quickly detect, contain and remediate advanced threats on your network.

- IBM X-Force Exchange
- OpenIOC
- IOC Bucket
- MISP
- Mandiant's IOC Finder
- ESET IOC Repository
- TAXII
- Splunk SA-SPLICE
- CybOX
- GitHub (google/grr Rapid Response for remote live forensics)

Leveraging the power of IOCs, you can find the footprints attackers leave behind when they breach your security defenses. It's one of the most effective ways to put advanced tactics to work to help protect against advanced threats.

About X-Force

Advanced threats are everywhere. Help minimize your risk with insights from the experts at IBM.

The IBM X-Force research and development team studies and monitors the latest threat trends including vulnerabilities, exploits, active attacks, viruses and other malware, spam, phishing, and malicious web content. In addition to advising customers and the general public about emerging and critical threats, IBM X-Force also delivers security content to help protect IBM customers from these threats.

IBM Security Services: Protect your enterprise while reducing cost and complexity

From infrastructure, data and application protection to cloud and managed security services, IBM Security Services has the expertise to help safeguard your company's critical assets. We protect some of the most sophisticated networks in the world and employ some of the best minds in the business.

IBM offers services to help you optimize your security program, stop advanced threats, protect data and safeguard cloud and mobile. Should you experience an IT security breach, IBM Emergency Response Services can provide real-time on-site support, including intelligence gathering, containment, eradication, recovery and compliance management. IBM Active Threat Assessment consulting services can help you identify hidden but active cyber threats before serious damage occurs to your infrastructure or even your brand. IBM Incident Response Planning can help you structure a cyber-security incident response plan (CSIRP) that incorporates the right process, tools and resources you need to respond to and help reduce the impact of a cyber attack. With IBM Managed Security Services, you can take advantage of industry-leading tools, security intelligence and expertise that will help you improve your security posture—often at a fraction of the cost of in-house security resources.

Contributors

Producing the IBM X-Force Threat Intelligence Quarterly report is a dedicated collaboration across all of IBM. We would like to thank the following individuals for their attention and contribution to the publication of this report.

Contributor	Title
David M. McMillen	Senior Threat Researcher, IBM Security Services
Leslie Horacek	Manager, IBM X-Force Threat Response
Luca Pugliese	Security Consultant, IBM Emergency Response Services
Michelle Alvarez	Threat Researcher & Editor, IBM Security Services
Pamela Cobb	Worldwide Portfolio Marketing Manager, IBM X-Force and Threat Portfolio

For more information

To learn more about IBM X-Force, please visit:
ibm.com/security/xforce/

[1] "RansomWeb: emerging website threat that may outshine DDoS, data theft and defacements?" *High-Tech Bridge Security Research*, 28 January 2015. https://www.htbridge.com/blog/ransomweb_emerging_website_threat.html

[2] Erica Chickowski, "Top 15 Indicators of Compromise," *InformationWeek: Dark Reading*, 09 October 2013. http://www.darkreading.com/attacks-breaches/top-15-indicators-of-compromise/d/d-id/1140647

[3] Lucas Zaichkowsky, "Hunting for Indicators of Compromise, *RSA Conference 2013*, 28 February 2013. https://www.rsaconference.com/writable/presentations/file_upload/end-r31.pdf

[4] Mandiant IOC Finder, http://www.mandiant.com/resources/download/ioc-finder/

[5] Nigel Willson, "APT Detection Indicators – Part 1," *Nige the Security Guy*, 12 December 2013. https://nigesecurityguy.wordpress.com/2013/12/12/apt-detection-indicators-part-1/

Please Recycle

IBM Security
Route 100
Somers, NY 10589

Produced in the United States of America
November 2015

WGL03099-USEN-00