Data breach digest.

Scenarios from the field.

verizon

Table of contents: the usual suspects

Welcome to the field....3
Mapping the industries, patterns and scenarios....5
The human element 8
Scenario 1: Social engineering – the Hyper Click 10
Scenario 2: Financial pretexting – the Slick Willie 15
Scenario 3: Digital extortion – the Boss Hogg....18
Scenario 4: Insider threat – the Rotten Apple 22
Scenario 5: Partner misuse – the Busted Chain.... 26
Conduit devices 30
Scenario 6: USB infection – the Porta Bella....31
Scenario 7: Peripheral tampering – the Bad Tuna 35
Scenario 8: Hacktivist attack – the Dark Shadow 38
Scenario 9: Rogue connection – the Imperfect Stranger 43
Scenario 10: Logic switch – the Soup Sammich.... 46
Configuration exploitation 49
Scenario 11: SQL injection – the Snake Bite 50
Scenario 12: CMS compromise – the Roman Holiday.... 54
Scenario 13: Backdoor access – the Alley Cat 58
Scenario 14: DNS tunneling – the Rabbit Hole.... 62
Malicious software 65
Scenario 15: Data ransomware – the Catch 22.... 66
Scenario 16: Sophisticated malware – the Flea Flicker 70
Scenario 17: RAM scraping – the Leaky Boot....74
Scenario 18: Credential theft – the Poached Egg 78
Way forward 82
Appendix A: Top 25 VERIS threat actions.... 83
Appendix B: CIS critical security controls 84

Welcome to the field

The Verizon RISK Team performs cyber investigations for hundreds of commercial enterprises and government agencies annually across the globe. In 2015, we were retained to investigate more than 500 cybersecurity incidents occurring in over 40 countries. In 2008, the results of our field investigations were the genesis of the first Data Breach Investigations Report (DBIR), an annual publication that dissects real-world data breaches with the goal of enlightening the public about the nature of the threat actors behind the attacks, the methods they use, including the data they seek, and the victims they target.

Our incident data corpus, which contains a wealth of breach information from over 70 contributors, has fueled eight years of DBIR reporting (to include this year's ninth DBIR). VERIS has allowed us to analyze this data to uncover the actors, actions, assets, and attributes involved in the incidents. While the DBIR focuses on trends and patterns found in an aggregated incident data set, the Data Breach Digest (DBD) gets you closer to the action. With the DBD, we're leveraging VERIS like never before. We're lashing up the cold, hard VERIS metrics to our on-the-ground casework experience. Essentially, we have opened our case files, and are giving you a first-hand look at cyber investigations from our experiences—a view from the field.

In 2015, we were retained to investigate more than 500 cybersecurity incidents occurring in over 40 countries.

A brief refresher on VERIS
VERIS, the Vocabulary for Event Recording and Incident Sharing, provides a common language for describing security incidents in a structured and repeatable manner. To facilitate the tracking and sharing of security incidents, we released VERIS for free public use. Get additional information on the VERIS community site;[1] the full schema is available on GitHub.[2] Both are good companion references to this compendium.

Many data breach victims believe they are in isolation, dealing with sophisticated tactics and zero-day malware never seen before—we see otherwise. To us, few breaches are unique. In fact, our VERIS research indicates that at any given point in time, a small number of breach scenarios comprise the vast majority of incidents we investigate. There is tremendous commonality in real-world cyber-attacks. In fact, according to our RISK Team incident data set over the previous three years, just 12 scenarios represent over 60% of our investigations.

1 veriscommunity.net/
2 github.com/vz-risk/veris

This is our opportunity to slice through the fear, uncertainty, and doubt that's so prevalent in security to reveal what's really happening in the cyber investigation field. These scenarios paint the picture behind the color-by-numbers VERIS data—they illustrate how breaches work, and include intrusion vectors, threat actions, and targeted vulnerabilities. Most of all, they help to prescribe a recipe for prevention, mitigation, and, if necessary, efficient and effective incident response.

For the DBD, we selected 18 data breach scenarios for two reasons: their prevalence and/or their lethality. In terms of prevalence (commonality), 12 scenarios met this criterion, while six scenarios met the lethality criterion. For lethality, at least two of the following three conditions had to occur: 1) difficulty in detection or containment; 2) level of sophistication; and/or 3) amount of potential resultant damage.

For the DBD, we selected 18 data breach scenarios for two reasons: their prevalence and/or their lethality. In terms of prevalence (commonality), 12 scenarios met this criterion; six scenarios met the lethality criterion.

Enemy courses of action (COAs)
Focusing on the "most prevalent" and to an extent the "most lethal" data breach scenarios is akin to the U.S. Army's approach for tactical field units preparing for combat. Within the "Five Paragraph" Operations Order, tactical units prepare for two possible enemy COAs: the "Most Likely COA" and the "Most Dangerous COA."

All scenarios draw from real-world cyber investigations. To protect victim anonymity, we modified certain details, taking some creative license. This included, but was not limited to, changing names, geographic locations, quantity of records stolen and monetary loss details.

The individual scenarios are the closest we can get to giving you a "RISK Team ride along." We hope you find these scenarios both interesting and informative. As you flip through the DBD, consider your level of exposure to similar threat actors and attacks. More specifically, ask yourself whether you have the Attack-Defend Card countermeasures in place, and what you can take away from the "Lessons learned" sections before it's too late. And, should that fail, how quickly you could enact the guidance given in the "Remediation and recovery" section.

If you find yourself responding to cybersecurity incidents, these scenarios shed some light on sources of evidence and methods allowing an investigation to progress quickly to containment and recovery (and, of course, improvement from the pre-incident state). Understanding the data breach scenarios that are relevant to you, your industry and your asset landscape is key to smart security. If you like what you find, we suggest you read the scenarios that affect other industries as well. Recognizing what helped or hindered investigative efforts in these cases will help you in future responses to cybersecurity incidents.

Understanding the data breach scenarios that are relevant to you, your industry and your asset landscape is key to smart security.

Enjoy the ride.

Mapping the industries, patterns and scenarios

Veteran DBIR readers might be asking themselves what the difference is between the data breach scenarios and incident classification patterns? A solid question. We will answer this after a brief introduction to incident classification patterns for non-veteran DBIR readers.

Incident classification patterns
To take this from tribal knowledge and a dash of data science, we utilized a statistical clustering technique that identified strongly related incidents and classified them into the nine incident classification patterns (this encompassed over 90% of our data corpus).

The incident classification patterns involving confirmed data breaches, in order of frequency, over the past three years are:

1. Point-of-sale (POS) intrusions – POS application/system related attacks.
2. Web app attacks – web application related stolen credentials or vulnerability exploits.
3. Cyberespionage – state-affiliated, targeted attacks.
4. Crimeware – malware used to compromise systems.
5. Insider and privilege misuse – unauthorized insider related activity.
6. Payment card skimmers – physically installed malicious card readers.
7. Miscellaneous errors – any mistake that compromises security.
8. Physical theft and loss – physical loss or theft of data/IT related assets.
9. Denial of service (DoS) attacks – non-breach related attacks affecting business operations.

Of the nine patterns, we chose the first six for this publication. Three were not included for the following reasons: 7 – Miscellaneous errors (mistakes, boring), 8 – Physical theft and loss (physical threat, not usually investigated by us), and 9 – DoS attacks (not a data breach).

Data breach scenarios
We clustered the 18 data breach scenarios into four groups, each described as follows:

A. The human element—five scenarios highlighting human threats or targets.
B. Conduit devices—five scenarios covering device misuse or tampering.
C. Configuration exploitation—four scenarios focusing on reconfigured or misconfigured settings.
D. Malicious software—four scenarios centering on sophisticated or special-purpose illicit software.

The key difference between "scenario" in this compendium and "pattern" is that the data breach scenarios are examples of specific incidents that fall into one of the six patterns.

Industries
We used the North American Industry Classification System (NAICS) for coding the targeted victim industries.[3]

Now let's use this thing!
The figure below provides the specs for the NAICS industries (left column), the six incident classification patterns covered in this publication (top row), and the corresponding most-relevant and semi-relevant scenarios (bottom rows).

The percentages below are based on VERIS metrics over the previous year: the "gold" boxes are those above 10%, the "red" boxes are those above 20%.

Incident pattern ►	POS intrusions	Web app attacks	Cyber-espionage	Crimeware	Insider and privilege misures	Payment card skimmers
Industry (NAICS #) ▼						
Accommodation (72)	53%	1%			3%	
Administrative (56)		4%	1%		6%	
Educational services (61)		9%	12%	22%	11%	
Entertainment (71)	58%	11%	11%		5%	
Financial services (52)		17%	1%	21%	6%	7%
Healthcare (62)	7%	8%	3%	3%	20%	
Information (51)		26%	9%	46%	1%	
Manufacturing (31-33)		3%	36%	19%	3%	
Mining (21)			11%		67%	6%
Other services (81)	1%	28%	3%	36%	6%	
Professional services (54)	2%	2%	26%	10%	1%	
Public (92)				18%	26%	
Retail (44-45)	20%	2%		25%	1%	4%
Transportation (48-49)			41%	9%	18%	5%
Utilities (22)		17%	50%	17%		
	▼	▼	▼	▼	▼	▼
Most-relevant scenarios – *the scenarios to read first!*	7, 10, 11, 12, 13 , 14, 16, 17, 18	8, 10, 11, 12, 14, 15, 16, 18	1, 6, 13, 14, 16, 18	2, 9, 13, 14, 15, 16, 18	3, 4, 5, 6, 13, 14, 16, 18	5, 7

Legend < 10% 11-19% > 20%

3 http://www.census.gov/eos/www/naics/

To use this publication, reference the chart above and follow these steps:

1. Select industry in left column.
2. Select incident classification pattern in top row.
3. Follow column down to bottom rows to identify scenarios by number relevant to chosen industry and enjoy the read!

For each data breach scenario, we provide an "Attack-Defend Card" along with a detailed situation. Data content within the Attack-Defend Cards is specific to the scenario (e.g., Social engineering – the Hyper Click) and is drawn from our RISK Team data set over the previous three years (unless otherwise specified). If applicable, scenarios that are considered as "lethal" are labeled as such. The contents are described as follows:

See Appendix A for a list of the "Top 25 VERIS Threat Actions". See Appendix B for a list of the "CIS Critical Security Controls."

The detailed situations associated with the data breach scenarios are drawn from our previous casework. Each situation walks you through initial detection and validation, response and investigation, and remediation and recovery.

While you peruse this compendium, take note of the scenario Attack-Defend Cards; assemble them into an Attack-Defend Card Battle Deck for use in your data breach prevention, mitigation and response preparedness efforts.

The human element

Leveraging human beings to gain access to information is not new; it predates binary. In our entire corpus of data breaches, we witness social tactics being used in around 20% of confirmed data breaches, only ranking behind the VERIS threat action categories of hacking and malware in prevalence. When looking only at the previous three years, the frequency increases to almost 30% of data breaches. While there are many tactics that can be unleashed to manipulate people, the top three, phishing (72%), pretexting (16%), and bribery/solicitation (10%), represent the vast majority of social actions in the real world.

We witness social tactics being used in around 20% of confirmed data breaches.

As one would expect, email is the primary means of communication to the target (72%) followed by in-person deception (18%) and phone calls (12%), with a small amount of overlap across the three means of communication. Social actions are typically part of a blended attack, with malware also present in 85% of data breaches and hacking found in 50% involving the human element.

While scenarios 1–3 focus on human beings as the targets of attack, scenarios 4–5 focus on people in trusted roles as the threat actors. With regard to the latter, 9% of confirmed data breaches over the previous three years were categorized in the insider and privilege misuse pattern. These were the top industries affected by social actions and privilege misuse and contributing to a data breach (previous three years):

- Social actions: financial services, manufacturing, professional services, public
- Insider and privilege misuse: financial services, accommodation, healthcare, public

Your employees and your business partners can be potential threat actors or targeted victims. It is important to not lose sight of the role humans play in data breaches.

Threat actors
The VERIS Framework categorizes threat actors as "external," "internal," and "partner." These three types of threat actors are described as follows:

- External threats – originate from sources outside of the organization and its network of partners. Examples include criminal groups, lone hackers, former employees and government entities. Typically, no trust or privilege is implied for external entities.
- Internal threats – are those originating from within the organization. This encompasses company full-time employees, independent contractors, interns and other staff. Insiders are trusted and privileged (some more than others).
- Partners – include any third party sharing a business relationship with the organization. This includes suppliers, vendors, hosting providers, outsourced IT support, etc. Some level of trust and privilege is usually implied between business partners.

Attack-Defend Card

Scenario 1.
Social engineering—the Hyper Click.

Data breach scenario

Frequency:* 16%

Sophistication level:

○ ○ 3 4 5

Composition:
Organized crime, state-affiliated

Incident pattern

Pattern:
Cyber-espionage

Time to discovery:

○ ○ W M ○

Time to containment:

○ ○ ○ M ○

Threat actor

Motive:
Financial, espionage

Disposition:
China, Argentina, North Korea, Russian Federation

Tactics and techniques:
Phishing, pretexting, backdoor, export data, spyware/keylogger, downloader, c2, capture stored data, use of stolen credentials

Targeted victim

Industries:
Manufacturing, professional services, public, information, utilities

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-2, CSC-3, CSC-8, CSC-12, CSC-14, CSC-17

Description
Social engineering attacks rely on influencing or tricking people into disclosing information or conducting an action, such as clicking on a hyperlink or opening an email attachment. Tactics may include deception, manipulation, pretexting, phishing, and other types of scams. Social engineering can be merely a part of a threat actor's overall methodology or the end game itself.

* + social

Upstream phishing for a downstream profit.

"It's the [hu]man, not the machine."—Chuck Yeager

Detection and validation
In one particular instance, a customer contacted the RISK Team with an issue involving a primary competitor, a suspected threat actor, located on another continent that had recently made public a new piece of large construction equipment. At first glance, the equipment looked like an exact copy of a model recently developed by our customer, the victim. This was even more suspicious as the competitor, the threat actor, hadn't traditionally produced this type of equipment and therefore had no past track record in this part of the market. The victim's concern was not just that this equipment's design details were obtained illicitly, but that other projects were also in danger of similar compromise.

Investigative approach
In data breach investigations, the response doesn't always only involve the analysis of digital evidence. In many of our cases, we find that traditional investigative techniques are just as important as, if not more so, than data obtained from the latest forensic tools.

In this case, interviewing the chief design engineer proved integral in determining how the design had been taken. By interviewing key employees, we were able to focus on the system used by the chief design engineer for the specific model of equipment that had possibly been stolen.

Response and investigation
Shortly after initial notification, we arrived onsite at the victim's headquarters and set about interviewing the key stakeholders. We began by working with the design team responsible for the equipment model that was the focus of the cyber investigation. In comparing features listed by the threat actor on their recently released model, the victim's design team identified several key parts and details that appeared identical to their own model. Many of these design elements were new and unique to the industry.

After determining that it was most likely that the equipment model designs had been compromised, our first request was for the names of those employees who worked on the design project for the equipment model involved in the design plan theft.

The first employee we interviewed was the chief design engineer for the project. While interviewing him, it became clear that he was actively looking for employment elsewhere and he might not be employed by the victim much longer. A recruiter had contacted the engineer via LinkedIn, which led to them exchanging emails.

A digital forensic examination of the chief design engineer's system and associated firewall logs provided evidence of a breach associated with the design plans, which were resident on that system. A PHP (scripting language) backdoor shell was found on the system. There were also clear indications that the threat actors had located and copied the file containing the design plans.

Malware spotlight: command and control (C2)
C2 refers to the methods or resources used by malware to communicate with its operators. C2 servers may be used to manage thousands of infected systems, and by issuing a single command from this system, they can all be marshalled into action. Advanced threats typically encrypt their C2 channels via the Secure Sockets Layer (SSL) encryption that is used in HTTPS or Secure Shell (SSH) connections. This encryption not only makes it harder for monitoring and detection solutions, but also makes it significantly harder to identify specific commands when C2 traffic is found.

In examining the engineer's email files, we found one from the recruiter occurring just prior to the beaconing activity. We then found an employment position-listing document attached to the email embedded with a small piece of malicious software (malware). Analysis of the malware revealed it contained a known-malicious Chinese IP address hard-coded within.

The stolen data included design blueprints for a new and innovative piece of large construction equipment. Through attack profiling, it was determined that the likely threat actors were a Chinese hacking group that had long been suspected of being state funded. Intelligence sources indicated that these threat actors had performed similar attacks against a variety of victims and allegedly provided the stolen intellectual property to Chinese companies that were state owned, operated or supported.

Through attack profiling, it was determined that the likely threat actors were a Chinese hacking group that had long been suspected of being state funded.

The threat actors had done their homework, as they identified the one key employee who would likely have access to the data they wanted—the chief design engineer for the project. The threat actors then established contact with the engineer through a LinkedIn profile under the guise of a recruiter with attractive employment positions and began sending emails containing fictitious employment opportunities. One of those emails contained an attachment that had a malware file embedded in the document. When opened, the malware began beaconing to an external IP address used by the threat actor. The threat actors then installed a backdoor PHP reverse shell on the chief design engineer's system.

From that beachhead, the threat actors were able to search the data on that system as well as collect sensitive data from network file servers and attached USB hard disk drives. At initial glance, the activity would almost seem normal, as the chief design engineer had legitimate access to all these data repositories. As he was deeply involved with this project, it wouldn't be suspicious for him to be accessing the various project-related files.

Upon completion of the data aggregation, the threat actors encrypted and compressed the intellectual property, and in doing so, made it unidentifiable to Data Loss Prevention (DLP). At that point, exfiltration was trivial and accomplished through an outbound HTTP connection. Unfortunately for the victim, the investigation confirmed that it had indeed lost intellectual property. Its suspicion that a foreign competitor leveraged the data in order to begin marketing a remarkably similar piece of equipment was substantiated.

Figure 1

1. Intial compromise via social media
2. C2 beaconing and commands
3. Attacker pivots to steal sensitive data
4. Data encrypted and exfiltrated

Remediation and recovery

With the chain of events clearly laid out, the victim then turned toward remediation. There was nothing it could do to recover the lost intellectual investment, but this victim was sure it did not want to go through this a second time. In many cases, this victim had done the right thing, but had still been breached. Especially with social threats, we find that even the most mature organizations can fall victim to data theft. We provided many recommendations, ranging from easy wins to more robust and involved solutions, which the victim worked into its current security posture.

One of our first recommendations was for the victim to set up a more comprehensive training and awareness program related to social engineering threats that employees may face. This focused on specific areas of the business and the types of information that were most critical to each job role. Clear steps were put in place to specify when and how data could be transferred. Part of this process was identifying information, such as new design plans, that should have additional security controls for proper handling. Engineers were provided with dedicated systems for them to perform their engineering work on, which no longer had email or web access. This would limit the number of avenues that potential threat actors would have to load malware onto these sensitive machines.

Social threats are hard to defend against, even when a good plan is in place, so we also recommended the victim adopt more robust monitoring solutions to identify the early signs of a compromise. Many of the core pieces of security existed – anti-virus deployments, intrusion detection sensors and NetFlow capture were all available, but mostly unused. Anti-virus was installed on all corporate assets, but the software was a mishmash of vendors as IT staff tastes changed over the years. We recommended selecting a single vendor and using a centralized solution so that updates could be rolled out across the company. Intrusion detection alerts and NetFlow capture can be correlated in many security event frameworks, and we suggested the victim take its existing infrastructure and centralize the results. Paired with the centralized anti-virus, these tools would allow IT and security teams to more quickly identify emergent threat actors before significant damage occurred.

Intrusion detection alerts and NetFlow capture can be correlated in many security event frameworks.

Lessons learned

Some of the measures an organization can take to reduce the impact of social engineering attacks may include a comprehensive and clear information security policy, user education through training and awareness programs and periodic audits to check policy compliance. Security controls can be enhanced with strong and mutual authentication combined with a robust identity and access management program.

Attack-Defend Card

Scenario 2. Financial pretexting—the Slick Willie.

Data breach scenario

Frequency:* 7%

Sophistication level: 2 3

Composition:
Organized crime

Incident pattern

Pattern:
Everything else

Time to discovery: W M

Time to containment: M

Threat actor

Motive:
Financial

Disposition:
Varies

Tactics and techniques:
Pretexting, privilege abuse, bribery, use of stolen credentials

Targeted victim

Industries:
Financial services, accommodation, retail

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-14, CSC-17

Description
Financial pretexting is a form of social engineering characterized by threat actors using false pretenses to trick or dupe a victim into performing a financial transaction or providing privileged data. These attacks use multiple communication channels and often employ social media networks. Typical threat actors are organized crime groups with the motive of financial gain.

* + social AND + motive.Financial

Almighty Zeus! Or not.

"You may fetter my leg, but Zeus himself cannot get the better of my free will."—Epictetus

Detection and validation
A regional banking organization contacted the RISK Team after a recommendation by its cyber insurance carrier. During an initial call with them, we soon learned that an unknown threat actor had attempted to initiate several wire transfers through the FedWire system totaling $5.3 million. Oddly enough, the bank didn't discover this attack; instead, they were notified by the Fed. The sheer dollar value of the transfers triggered a volumetric alert from the Fed once the bank's own required reserves dropped under $500K. Luckily, because of this no transfers were successful.

Cyber insurance carriers: covering your digital assets
Last year, cybersecurity incidents cost companies hundreds of billions of dollars. These costs included everything from public relations and crisis management consulting, forensic investigation costs, outside legal counsel, credit monitoring, notifications, call centers and more. In addition to monetary loss, breached companies faced other, more intangible "losses," such as business interruption, reputational damage, litigation and regulatory actions. Generally, cybersecurity related risks aren't covered as part of traditional insurance policies. Enter cyber insurance carriers.

Although not necessarily a new insurance offering, cyber insurance has become more and more a factor in incident response (as well as disaster recovery) planning. Cyber insurance coverage is used for cybersecurity incidents involving data loss and destruction, DDoS attacks, malware outbreaks, etc. Beyond first-party loss and third party indemnification, this coverage may also include other benefits such as public relations support, investigative response expenses and security audits.

As to be expected, no insurance carrier is going to blindly accept unlimited quantities of risk. They too have to manage their risk and in turn will typically have specific requirements of the insured. While these requirements for coverage vary by carrier, the basic premise is that the insured must have and maintain an adequate IT security program. As time has gone on, cyber insurance carriers have played an ever increasing role in driving up security awareness and driving down risk. As this evolution continues, more insurers are adding enhanced customer benefits and/or discounts on premiums based on the insured's cybersecurity maturity.

In addition to working with an ever increasing number of cyber insurance carriers as it relates to specific data breach investigations, we're also finding and encouraging a closer proactive security relationship between the insured and their cyber insurance carriers—that is, joint breach simulations, table-top exercises, and policy review and development.

Response and investigation
During the initial interviews with the victim, we learned that a manager in its finance department had initiated requests for multiple wire transfers over a 24-hour period using the bank's FedWire application. We interviewed the finance manager and found that she was completely unaware of the attempted transfers. She indicated that over the previous couple of weeks her system had been "acting funny." She then commented that it "does things on its own sometimes." To investigators, a statement like that translates to "my system has been completely pwned."

Earlier that month the finance manager had received an email purportedly from the bank's CIO. She told us that the CIO had sent her a glowing message stating that members of his team had mentioned what a great business partner the finance manager had been on recent collaborative projects. The message went on to state that the CIO wanted to make sure he personally recognized employees that were as highly regarded as the finance manager since they were such an asset to the business. The message contained what appeared to be an innocuous hyperlink, which the finance manager recalled clicking. She thought this email was odd given that to the best of her recollection she had never worked with the bank's CIO or any of his team members.

Reading the metadata tea leaves
Analysis of the metadata in the email revealed that it had been spoofed to appear to be a legitimate company email. In fact, the email used the CIO's full first name, middle initial and last name, whereas the legitimate company email had only the first and last name of the CIO. Had the victim created filters to spot and block this type of attack, this form of malicious email could have been easily identified as spoofed by the email gateway.

Upon speaking to the CIO, it became clear that he didn't even recognize the finance manager's name; much less had he sent her a "pat on the back" email. We then ran an anti-virus scan against an image of the finance manager's computer system – and again, to no one's surprise, we found a Zeus Trojan infection dating back to when the CIO email was received.

Remediation and recovery
This Zeus variant was particularly nasty in that beyond the standard credential-stealing and data-scraping capabilities of a standard infection, it allowed for full remote access and control of the affected system. Presumably, this would allow a threat actor to initiate fraudulent wire transfers from the appropriate system inside an organization. Not only would the threat actor have the appropriate credentials to pull off the crime, it would be doing it from the right system.

The hyperlink in the fake email connected to a Zeus installer, which was still an active host when we accessed the email. The message itself was almost perfectly targeted. It was sent to one of only two employees at the branch who was authorized to initiate FedWire wire transfers and was written in such a way as to adequately put her in the correct frame of mind before asking her to click on a malicious hyperlink.

The simple fact is that this crime would likely have worked if the threat actors had kept their greed in check. As noted, the transfers were stopped based on volumetric alerting; essentially, the threat actors tried to move too much money.

Lessons learned
Threat actors who engage in social engineering attacks do it because they know that the human element is the weakest link in any information security strategy. They often take advantage of their targeted victim's sense of curiosity and psychology in order to gain access to sensitive data. In this particular case, the victim employee was showered with compliments by a company executive and then asked to click an innocent looking hyperlink. Bottom line: employees need to be constantly sensitized and trained through security awareness programs in order to be extra vigilant regarding their actions. Multi-factor authentication should be implemented wherever feasible for access to financial systems to combat reuse of stolen credentials.

Threat actors who engage in social engineering attacks do it because they know that the human element is the weakest link in any information security strategy.

Attack-Defend Card

Scenario 3 [Lethal]. Digital extortion—the Boss Hogg.

Data breach scenario

Frequency:* 9%

Sophistication level: 2

Composition:
Organized crime

Incident pattern

Pattern:
Everything else

Time to discovery: H D

Time to containment: D

Threat actor

Motive:
Financial

Disposition:
Varies

Tactics and techniques:
Extortion, ransomware

Targeted victim

Industries:
Financial services, public

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-8, CSC-10

Description
Extortion cases are often difficult and harrowing for the victim. As investigators work furiously to validate claims made by the threat actor, the doomsday clock continues to tick. Often, victims are faced with the inevitable decision to give in, or not give in, to threat actor demands as the clock counts down. For this reason, we consider digital extortion as a lethal data breach scenario.

* + social (extortion) OR malware (ransomware) including all incidents, not just data breaches

The Shakedown, takedown.

"Pay da man his money. He earned it straight up."—Teddy KGB

Incident handling focus: extortion demands
Extortion is nothing new to the rich and powerful and it was only a matter of time before this age-old tactic was digitized. Digital extortionists have a wide variety of targets—an individual's personal information, such as pictures or documents, as well as a company's secret files or customer information. The effects of such an attack are, at best, frustrating and have the potential to result in serious data loss or operational impacts.

We typically recommend that victims never acquiesce to extortion demands. As is the case with many negotiation scenarios, giving in to extortionist demands simply adds fuel to the fire. Instead, we recommend that organizations and individuals take steps prior to an incident to mitigate risks. Backup solutions to recover "ransomed data" can help to defuse many extortion demands. Customer data is often a target, especially for those seeking to damage a corporate reputation. Plans for how to handle the situation where data is being held hostage should be created to minimize response times and provide a clear path of action.

Detection and validation
The RISK Team was contacted by a large-scale manufacturer and retailer of consumer goods in North America. A member of the IT infrastructure team at this particular organization received two separate emails from an individual in Southeast Asia claiming to have successfully exfiltrated several years' worth of customer order data. The individual was seeking monetary payment in exchange for not releasing the data publicly. The first of these two emails seemed innocuous, almost like a simple spam message. As such, the recipient simply chose not to respond. The second email took on a much more serious tone. It not only demanded $50K to withhold release of the data, but it also included a data sample to add legitimacy to the claim.

When the recipient verified the contents of the second email, he immediately notified the IT security team, who, in turn triggered the Verizon Rapid Response Retainer service. We received the call and were onsite the following day; our task: proving or disproving the claims made by the extortionist.

Our initial focus was two-fold: 1) determine how the data was stolen and 2) determine the scope of the theft. Based on the content of the sample data, we were able to direct our attention to the victim's e-commerce platform. Extortion demands aside, it was also imperative that we find the vulnerability that could be leveraged by future threat actors.

Our initial focus was two-fold: 1) determine how the data was stolen and 2) determine the scope of the theft.

Response and investigation
We started with a basic review of the e-commerce platform for obvious vulnerabilities and within a few hours discovered a weakness in the application's authentication mechanism. This vulnerability provided a threat actor with the ability to "force browse"[4] purchase confirmation pages and, in turn, view transaction details associated with customer purchases. This was accomplished by altering the URL string of any standard purchase confirmation page. A threat actor could access transaction details of essentially any transaction existing within the victim's database simply by changing the order number included in the URL string.

4 https://www.owasp.org/index.php/Forced_browsing

After validating the vulnerability, we reviewed the access logs on the e-commerce servers. We confirmed that the threat actor did use it to illicitly access the mother lode of customer data—essentially, several hundred gigabytes of HTML-based transaction information. During a roughly four-week timeframe, the threat actor ran a script that accessed the back-end database driving the victim's e-commerce platform, and exfiltrated over 1.5 million customer orders. Upon looking at the associated web server logs, the attack jumped out at us, plain as day, with a single source viewing absurd amounts of sequentially listed customer orders.

Remediation and recovery

With the attack vector and data impact confirmed, the victim refocused on responding to the threat actor. It's intent was to cut off the threat actor at the knees. The victim had decided that there was simply no way it was going to pay up; instead, it took away the only leverage the threat actor was counting on: the public shock value associated with the impending data release.

Rather than letting the threat actor release the stolen data, the victim beat them to it. The victim marshaled its public affairs team, then went public and fully disclosed that it had been breached. With a contrite and steady resolve, it admitted to the world, its customers and its shareholders, that it had lost nearly two million customer records over what was ultimately a low-hanging-fruit vulnerability. The organization offered a sincere apology and pledged to do better.

The victim marshaled its public affairs team, then went public and fully disclosed that it had been breached.

After that, the victim completely dismantled its e-commerce architecture and started over from scratch. The victim rebuilt the environment from the ground up with a full testing and development procedure, to include routine vulnerability scanning and penetration testing—something that was never really done before.

Lessons learned

Although the information disclosed on the transaction pages was not easily monetizable (that is, there was no Payment Card Industry (PCI) or banking data), financially motivated threat actors still conducted opportunistic attacks, targeting specific weaknesses exposed to the internet that were easily discoverable and exploiting them using automation. The victim incorporated non-technical business units (for example, public relations, legal, etc.) into its overall response to the incident and disclosed the data breach on its own terms. In the end, this victim took complete ownership of its data breach experience and used it as an opportunity to come out stronger and more secure on the other side.

Leveraging attacker hubris
On more than one occasion, the RISK Team has collaborated with law enforcement to capture cyber extortionists using a very basic technique – we offer them jobs. It goes like this...

In one case, we coordinated with law enforcement to fly an extortionist from Eastern Europe to Washington, D.C., for a job interview after he'd successfully gutted a financial services firm of several years' of customer records. Posing as the company's executives, we conducted the interview. After 30 minutes of regular interview questions, we posed the one we had been saving until last – "We're interested in hiring you. But, we need to know that it truly was you who accessed our systems. We need to know that you didn't just hire somebody else to do it. So, from the very beginning, tell us how you gained access to our systems."

The threat actor went on to describe how he used Structured Query Language (SQL) injection against the corporate website to open a remote command shell to the corporate web server. From there he was able to install malware capable of enumerating system-level administrator credentials and navigate through various protocols (mainly Remote Desktop Protocol (RDP), port 3389) to reach every system in the environment. The network was completely flat, he explained (and he was right, it was), which allowed him to jump between different business units (HR, Payroll, etc.) without traversing any kind of legitimate firewall. He explained that he then compiled as many interesting and sensitive documents as he could find and simply FTP'd them to himself from an end user system.

After taking the time to explain in excruciating detail the entire timeline of his technical achievements, the threat actor looked up and smugly asked, "So, do I have the job or what?"

To which our law enforcement partners replied, "No. No you don't. But thank you for the confession!"

Attack-Defend Card

Scenario 4.
Insider threat—the Rotten Apple.

Data breach scenario

Frequency:* 12%

Sophistication level: 1

Composition:
Cashier/bank teller/waiter, end users, organized crime, finance employees, call center employees

Incident pattern

Pattern:
Insider and privilege misuse

Time to discovery: D W M

Time to containment: D W M

Threat actor

Motive:
Financial, espionage, grudge

Disposition:
Varies

Tactics and techniques:
Misuse of physical and logical access, use of unapproved hardware, bribery

Targeted victim

Industries:
Financial services, accommodation, healthcare, public

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-5, CSC-6, CSC-13, CSC-16

Description
Although insider threat actors tend to be less frequent than their external counterparts are, we readily admit they make up some of the most interesting cases. Insider related data breaches involve threat actors with some level of trust and privilege causing a data breach through malicious intent.

* + misuse

Special "privileged" abuse.

"The greater the power, the more dangerous the abuse."
—Edmund Burke

Detection and validation
The RISK Team was called in to investigate an insider threat-related data breach. An organization was in the middle of a buyout and was utilizing retention contracts to prevent employee attrition. Based on an anonymous tip from an employee, suspicion was raised that a middle manager, hereafter referred to as "John," had access to, and was abusing, the CEO's email account.

Response and investigation
Late one evening after the employees had left the building, we arrived to meet with the Director of IT. He had no knowledge – nor the apparent "need to know" – of the incident, but was there to provide us with access to the systems and data. We worked throughout the night to perform forensic acquisitions of the CEO's system, the suspect's system, web-based email logs, and sundry other evidence sources. At just past midnight, we finally received the access we needed and were ready to dig-deeper, as our IT contact took off for home in search of some zzzs.

We needed to quickly establish if there was any truth to the claim that the middle manager was reading the CEO's email. Was it possible that the CEO's email archive was being shared across the network? Did the suspect have access rights to the CEO's mailbox through Microsoft Exchange? Was the suspect accessing the CEO's email through Microsoft Outlook Web Access (OWA)? The answer to all these questions was ultimately "no." While there are many ways to view someone's email, our cursory review of the system images and associated logs yielded nothing.

As the next day drew on, the lack of a "smoking gun," not to mention sleep, left our brains fried. After hitting the vending machine, we refocused and changed our approach. We swung back to the basics, started brainstorming, and sharpened Occam's razor by asking ourselves the simplest questions: How does email come into an organization? It usually comes from the internet through some spam filter before hitting the mail server. Did this organization have an onsite spam filter? Yes, a quick glance at a crude network diagram showed a standard spam filter setup.

We swung back to the basics, started brainstorming, and sharpened Occam's razor by asking ourselves the simplest questions.

The appliance itself wasn't a standardized system that we could acquire forensically. With credentials provided by our IT contact, we logged in and noticed that the filter was set up to log all incoming emails including the CEO's. This was a bit odd, but not necessarily unusual. A speedy check for the access logs to this appliance revealed that they had been recently deleted. We felt like we were onto something.

At this point, we needed to know who had access to the spam filter. Apparently, a few IT administrators had access, and none of them was John. In casual conversations with the IT director, we inquired about personal relationships between John and the short list of other employees. Bingo! It just so happened that one of the IT administrators, hereafter referred to as "Kevin," was very good friends with John.

Armed with this nugget of knowledge, we took an image of Kevin's system. Like John's, Kevin's system had zero in terms of web-browsing history. Thanks to our insight gained from the spam filter, we knew exactly which text "strings" to look for. A keyword search of the unallocated clusters (currently unused space potentially containing artifacts of previous activity) on both systems revealed strings

associated with logging into the spam filer and looking at the CEO's incoming email through good ole Kevin's administrator account. It turns out that Kevin had given John his credentials to log into the appliance and read incoming email for potentially any employee. In addition, John's system showed signs of having used Kevin's credentials to browse sensitive file shares and conduct other unauthorized actions.

"Ask the data"
A peek into the incident data that feeds into the DBIR shows that unlike this example, the majority (63%) of data breaches over the previous three years involving "insider and privilege misuse" were financially motivated. End-users with access to Personally Identifiable Information (PII) and bank employees with access to banking information are more prevalent than system administrators using privileged access. A pessimist would argue that this is because misuse leading to identity theft or fraudulent transactions is only identified as a result of the post-compromise fraud.

Remediation and recovery
We promptly reported our findings to the CEO, who then informed the legal and human resource (HR) departments. Soon thereafter, the decision was made to interview the two employees before moving forward. During the interviews, both employees denied any association with the spam filter, the CEO's email and the sensitive file shares. But the facts uncovered by our investigation left no doubt of the facts. After having worked a few insider cases, you begin to learn that most people, no matter how hard they try, or how comfortable they feel, aren't very good liars.

Upon completion of the interviews, the two employees in question received personal escorts out of the building. Needless to say, after this incident, the firm revisited its spam filter policy by reconfiguring it to log only flagged messages.

After having worked a few insider cases, you begin to learn that most people, no matter how hard they try, or how comfortable they feel, aren't very good liars.

"Bob, the force-multiplier"
One of the most memorable insider cases we have ever seen involved a US-based company asking for our help in understanding some anomalous activity that it was witnessing in its Virtual Private Network (VPN) logs. This organization had been slowly moving toward a more telecommuting-oriented workforce, and had therefore started to allow developers to work from home on certain days. In order to accomplish this, it had set up a fairly standard VPN concentrator approximately two years prior to this event.

The IT security department decided that it should start actively monitoring logs being generated at the VPN concentrator. It began scrutinizing daily VPN connections into its environment, and before long found an open and active VPN connection from Asia! When one considers that this company fell into the designation of US critical infrastructure, it's hard to overstate the possible implications of such an occurrence.

The company had implemented two-factor authentication for these VPN connections. The second factor was a rotating token key fob. The developer whose credentials were being used was sitting at his desk in the office. Plainly stated, the VPN logs showed him logged in from China, yet the employee was right there, sitting at his desk, staring into his monitor. The company initially suspected some kind of unknown malware that was able to route traffic from a trusted internal connection to China and then back. What other explanation could there be?

As it turns out, Bob had simply outsourced his own job to a foreign consulting firm. Bob spent less than one fifth of his six-figure salary paying a foreign firm to do his job for him. Authentication was no problem. He physically FedEx'd his token to Asia so that the third party contractor could login under his credentials during the workday. It appeared that Bob was working an average 9 to 5 workday. Investigators checked his web-browsing history, and that told the whole story.

A typical "work day" for Bob looked like this:

9:00 AM – Arrive and surf Reddit for a couple of hours. Watch cat videos.

11:30 AM – Take lunch.

1:00 PM – eBay time.

2:00ish PM – Facebook updates and LinkedIn.

4:30 PM – End of day update email to management.

5:00 PM – Go home.

Evidence even suggested he had the same scam going across multiple companies in the area. All told, it looked like he earned several hundred thousand dollars a year, and only had to pay the foreign consulting firm about $50K annually. The best part? Investigators had the opportunity to read through his performance reviews while working alongside HR. For the past several years in a row, he received excellent remarks. His code was clean, well written, and submitted in a timely fashion. Quarter after quarter, his performance review noted him as the best developer in the building. Nice work, Bob!

Attack-Defend Card

Scenario 5 [Lethal]. Partner misuse — the Busted Chain.

Data breach scenario

Frequency:* 4%

Sophistication level: 1

Composition:
Business-2-business partner

Incident pattern

Pattern:
Insider and privilege misuse

Time to discovery: D W M

Time to containment: D W M

Threat actor

Motive:
Financial, espionage

Disposition:
Varies

Tactics and techniques:
Misuse of physical and logical access, data mishandling

Targeted victim

Industries:
Financial services, accommodation, healthcare, public

Attributes:
Confidentiality

Countermeasures:
CSC-12, CSC-16

Description
Doing business today requires trusted relationships with business partners and vendors. Partners can manage critical devices, store or aggregate sensitive data, and/or be provided with remote access into corporate networks. Just as employees may have malicious intentions, vendors and business partners may also leverage legitimate logical or physical access for unsanctioned access to data. Because of this, we consider partner misuse as a lethal data breach scenario.

* + misuse AND + partner involving 'All' patterns

Late night fuel "benefits."

"Technologies tend to undermine community and encourage individualism."—Henry Mintzberg

Detection and validation
The RISK Team was retained by a cyber insurance carrier to investigate an unusual pattern of payment card fraud emanating from one of its customers—an oil and gas company. The insured operated a certain brand of gas stations, which we will refer to as the "Dixie Boys Truck Stop" (DBTS). We had an exploratory call with representatives of the insurer, the DBTS IT security team, and local law enforcement to discuss the background and facts of the case and agree on a path forward.

On the exploratory call, we learned that an escalating pattern of counterfeit fraud had begun at a single gas station about a month prior, then spread to five other locations. Analysis of the fraud patterns suggested the likely compromise of payment card magnetic-stripe data somewhere at DBTS, either at individual gas stations or at the headquarters. Bottom line: we needed to move quickly to identify the source of exposure and contain it.

Response and investigation
Given the information available, we set up shop at DBTS corporate headquarters and dug in. We initially focused on understanding the payment card transaction flow from the gas stations to the headquarters, and all the systems in between. Concurrent to mapping transaction flows, we were granted consent to cross correlate their internet communications against our cyber intelligence sources to identify connections with known threat actors or previously-conducted investigations. During the network traffic analysis, which yielded nothing suspicious, we discovered three additional gas stations were red flagged as demonstrating the same counterfeit fraud pattern. The ante had been upped and now included up to nine locations, with possibly more to follow. Pressure to contain the data breach was mounting fast.

Until we arrived at DBTS, several systems at each gas station were openly accessible from the Internet over certain high-numbered ports. All signs pointed to an external source for the data breach, although network and endpoint forensics conducted at the gas stations revealed none of the evidence we would expect to be present in a POS intrusion. All connections to these systems, both on the console and remote, could be accounted for and no indications of malware were present. Inspection of video camera footage from three stations showed no credible evidence of skimming at the cash registers or at the pumps. This ruled out the most common attack methods; obviously, something else was happening that was thus far undetected.

In coordination with law enforcement and with DBTS' support, we configured evidence traps on the payment processing servers at a number of gas stations within proximity to those showing fraud. The traps consisted of several moving parts, including keystroke logging, file integrity monitoring with alerting and playback recording of remote support sessions. Together with law enforcement, we immediately setup alerts. Within days, an alarm was tripped.

The traps consisted of several moving parts, including keystroke logging, file integrity monitoring with alerting and playback recording of remote support sessions.

We collected evidence; what these sources revealed was quite shocking. First, the support vendor, contracted by DBTS to provide general IT and POS support to the gas stations, connected via Remote Desktop over VPN to the payment processing server. Upon connection, a check occurred to verify no other active logons were in progress. Next, the system clock was set forward two years. Then, a configuration file was modified to enable a verbose debug setting in the payment application, creating an output file capturing clear text copies of authorization requests from

each fuel pump. This included complete mag-stripe sequences sufficient for conducting payment card fraud. The session ended with setting the clock back to the correct date and time.

The question then became whether the vendor was the source of the data breach or whether it was an upstream victim. The session was confirmed to originate from the vendor's support center, located nearby in the same city; the focus of the joint investigation then shifted to this location.

Upon notifying DBTS, we learned of no trouble tickets or outages reported at that gas station in the previous week. Law enforcement recommended we continue monitoring the evidence traps to detect attempts at harvesting the captured data.

The vendor as the vector

The data shows that you can't blame an investigator for focusing first on an external remote attack. Over the last three years of breach data, 99% of confirmed data breaches involving one or more asset varieties—to include POS terminal, POS controller, fuel pump terminal were driven by external threat actors.

Even after the source of the attack was identified as the vendor, it would be presumptuous to assume the vendor was driving the attack. We have seen numerous cases where POS vendors are targeted with phishing campaigns with the ultimate goal of stealing the credentials used to access their customer's POS environments. Using the same credentials across several clients makes this attack method even more successful. From a VERIS incident classification perspective, although partner credentials are being used, and potentially from their own compromised systems, we would still align this to an external threat actor, not a partner threat actor.

This drives home the point that not only should you secure remote access to critical environments by whitelisting source IP addresses, but two-factor authentication forces the threat actor to figure out another way in, and in doing so, raises the bar. While we still unfortunately analyze numerous POS smash-and-grab jobs, we have seen the level of effort begin to rise from attacking internet-facing POS systems with default credentials to leveraging stolen POS vendor credentials. We hope that eventually the target base will shrink significantly for threat actors relying on stolen single-factor credential reuse to access sensitive data.

There was very little time for thumb-twiddling as the next evening another alert was tripped from the same victim location, also originating from the vendor. This time, playback showed the system clock was set backward in time. Next, the debug file was opened, followed by a "ctrl-c," and then the file was closed. The system clock was set back to current and the session terminated. As planned, we immediately notified law enforcement, which dispatched a patrol car to check out the vendor's support center.

Remediation and recovery

Being a Saturday night, there was only a single car in the parking lot, which pointed to a particular member of the vendor's helpdesk staff as the most likely threat actor. Exactly what transpired at that moment between the officer and employee is unknown, but it clearly led to some sort of confession.

As it turned out, this individual would seek out late-night assignments over the weekends that required only a single person in the office on call. He would connect to customer systems to steal payment card data, believing that the time of day would provide less risk of being caught in the act, and attempted to cover his tracks with what were ultimately useless anti-forensic techniques. This particular

threat actor attempted to distance himself from the privilege misuse by conducting all malicious activity solely on his manager's desktop system.

Lessons learned

Several takeaways surfaced from this breach. DBTS' assumption that its POS vendor had implemented security practices was a procedural omission. The fact that the threat actor was able to use another system to attempt to cover his tracks shows that shared logins were used by the help desk staff to perform operations. The shared logins limited accountability and gave the threat actor the confidence that he could get away with it; in this case, he obviously didn't, but unique user account logins could have discouraged the attempt.

Two-factor authentication was not utilized for remote access into the POS servers. From a threat modeling perspective, a keylogger on any of the help desk systems is all that it would take for this to morph from a partner misuse breach to a widespread external breach.

Help, my payment card environment has been breached!

For merchants who signed a contract with their acquiring bank and at least one payment card brand, they have agreed to meet the requirements of the Payment Card Industry Data Security Standard (PCI DSS). PCI DSS requires that merchants experiencing a security issue defined as an actual or suspected compromise of payment card data notify their acquiring banks and affected payment card brands. Additionally, these victims may be required to engage expertise approved as part of the PCI Forensic Investigator (PFI) Program to investigate the security issue.

The PFI conducts an independent investigation of the incident to determine what happened, how it happened, and what specific payment cards were or may have been compromised. Each individual payment card brand provides input via the PCI Council to the PFI investigations process as it relates to their brand. The PFI has been through training, has experience in conducting these investigations and works with data breach victims to meet the requirements for the PFI program, as well those for each payment card brand.

As the investigation progresses, the PFI often makes recommendations relating to breach containment. In addition, the PFI makes long-term recommendations for improving security as it relates to payment card data. A final management report of the investigation, including findings, scope of the breach and recommendations is provided at the end of the investigation. Victims are required to provide this PFI report to their acquiring bank and any affected payment card brands.

Conduit devices

Devices play a significant role in data breaches as assets targeted either for the data they store/process or because of their accessibility to the outside world. We often see devices used as tools to advance an attack, as well as device ownership and management often appearing as a factor in data breaches.

Two DBIR incident classification patterns are constructed around the assets affected: POS intrusions and payment card skimmers. These two patterns account for 40% of the data breaches in our VERIS data set.

The top five asset varieties represented in confirmed data breaches are:

1. POS controllers – these process a highly sought after data variety (payment cards).
2. Desktops – these provide a common foothold into a corporate environment.
3. POS terminals – as with the POS controller (above), these process a highly sought-after data variety.
4. Web apps – many times the lone internet-facing device for an organization.
5. People – not commonly thought of as an asset, but they actually are when targeted in social attacks.

Attack-Defend Card

Scenario 6. USB infection—the Porta Bella.

Data breach scenario

Frequency:* (4% overall) 33%

Sophistication level: 4 5

Composition:
State-affiliated, organized crime

Incident pattern

Pattern:
Insider and privilege misuse

Time to discovery: D W M

Time to containment: D W M

Threat actor

Motive:
Financial, espionage

Disposition:
China, North Korea, Russian Federation

Tactics and techniques:
Use of unapproved hardware, pretexting

Targeted victim

Industries:
Manufacturing, professional services, public

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-3, CSC-7, CSC-8, CSC-12, CSC-13, CSC-17

Description
Digital denizens are familiar with the USB flash drives given away at trade shows, conferences, information booths and the like. Some are even sent to recipients via snail mail, pre-loaded with useful marketing data and pre-configured to auto-link to websites upon initiation. These handy devices are ubiquitous among swag bag collectibles along with ink pens that don't work, stale mints and badge lanyards of all descriptions. However, unlike the other conference detritus, these drives can carry a dangerous payload.

* + use of unapproved hardware

Gone with a flash!

"It is the small things in life which count; it is the inconsequential leak which empties the biggest reservoir."—Charles Comiskey

Detection and validation
After returning from an industry conference, a film industry executive received an envelope that looked like it was from a well-known production company. The envelope contained correspondence on company letterhead and a branded USB flash drive. The letter requested that the executive review the press kit contained on the drive. Recalling the production company's booth at the conference, the executive inserted the USB flash drive into his laptop system and opened an enticingly-named executable file.

"Ask the data"
The efforts taken to add legitimacy to the mailing (letterhead, branding on the USB flash drive) and the intelligence gathering to leverage the timing of the industry conference—not to mention the delivery method of snail mail addressed to a particular target—allow us to easily categorize this as a targeted attack.

Based on our incident corpus, these targeted attacks represent just over a third of confirmed data breaches associated with malware. If we include all malware incidents from our data set (that is, installation of malware was confirmed but data loss was not), the percentage goes down from approximately 33% to 16%. This makes sense if one considers the incident classification pattern of crimeware and the prevalence of malware that is served up on compromised sites or malicious advertisements. These "run of the mill" infections would be expected to be more commonplace than the ultra-sophisticated attacks

After the movie trailer finished, the executive closed the window and reviewed the rest of the press kit files. The executive thought nothing of it, because the files closely resembled those of a press kit that a production company usually releases as part of its promotions during the holiday movie season.

Response and investigation
Upon execution, the executable file did two things. First, it played a trailer for an upcoming movie from the production company and then it silently installed malware on the system with the aim of stealing an unreleased movie. The malware established persistence via Windows Registry key entries and attempted to reach out to a C2 server. Thankfully, the executive's company had a proxy server that monitored all outgoing traffic from the corporate network. The proxy server blocked an attempted connection to a C2 server and forwarded a low-level alert to the IT security team.

The blocked connection alert was one of many alerts that were forwarded to the security analyst each day. Since it blocked access to the blacklisted IP address, the analyst didn't chalk it up as a success and moved on. It wasn't until his manager initiated a full-fledged security review that he dug a little deeper. Our review of the logs found that while the proxy had blocked the initial connection attempt, it had allowed an encrypted connection to another server. While that server wasn't on the proxy's blacklist, it was suspicious.

Given the nature of the security review, the analyst initiated the company's Incident Response (IR) playbook and contacted the company's IR team. After the analyst explained the situation, the IR manager contacted the executive and asked him to recount the events leading up to the alert. After the discussion, the IR manager asked the executive to disconnect the system from the network and leave it powered on without internet access. In addition, the IR manager contacted the RISK Team. She wasn't certain that the USB flash drive was the culprit, but thought she had enough to get an investigation started.

Within 18 hours of the incident, we were on site conducting interviews with the executive, the analyst and the IR manager. After collecting the relevant logs, the USB flash drive, a physical memory dump and a forensic image of the system, we sent the evidence items to the RISK Labs for analysis. We also started collecting physical memory dumps and forensic images of other systems that had access to the final cut of the movie.

After collecting the relevant logs, the USB flash drive, a physical memory dump and a forensic image of the system, we sent the evidence items to the RISK Labs for analysis.

USB flash drive forensics
We were fortunate to have the malware from the USB flash drive and a forensic image of the system. Usually, a USB flash drive is long gone by the time we arrive onsite and we have to use endpoint forensics of the system to piece together the overall picture.

Some people may believe they are safe from USB flash drive attacks because their DLP software blocks said devices. Unfortunately, DLP can't block all USB devices without making systems unusable. For example, a USB device claiming to be a Keyboard Human Interface Device can bypass DLP and "type" scripted keystrokes at one-thousand words per minute. It acts just like a keyboard and basically gives keyboard access to the threat actor. With easy to use tools to load custom payloads on to the device, anything is possible.

Analysis of the malware on the USB flash drive and the laptop showed that it was capable of establishing a reverse shell. Once the reverse shell had been established, it didn't take long for the threat actor to start exfiltrating gigabytes of data including an unreleased movie.

Remediation and recovery
Because some of the files made it out onto torrent sites, containment and eradication began within a week of the incident. While the internal IR team removed the laptop from the network, it wasn't certain whether the malware had spread or whether unauthorized parties had accessed other confidential data. After the images of the laptop were collected, we advised the IR team that the laptop needed to be rebuilt.

Using the IOCs we provided in the report, the IR team confirmed that the laptop no longer contained the malware. Working in conjunction with the IR team and the IT security team, we found that the logs indicated the malware hadn't spread to any additional systems. To confirm, the IR team and IT security team scanned its network for the IOCs from the laptop and found nothing.

In our investigation report, we recommended that the company expand its access to security intelligence and use that intelligence to feed its security devices like the proxy server. Even the anti-virus solution failed to detect the malicious activity. We also recommended the company double-check that its endpoint anti-virus solutions are installed and running with the latest definitions. Of course, additional education regarding the proper use of USB flash drives and corporate assets was also necessary. Finally, we recommended that the user's credentials be reset on all corporate-owned assets and that they consider the same action for any users that worked closely with the executive during the time of compromise.

Attack-Defend Card

Scenario 7. Peripheral tampering—the Bad Tuna.

Data breach scenario

Frequency:* <1%

Sophistication level: 2

Composition:
Organized crime

Incident pattern

Pattern:
Payment card skimmers

Time to discovery: H D

Time to containment: M

Threat actor

Motive:
Financial

Disposition:
Bulgaria, Romania, Armenia, Brazil, the U.S.

Tactics and techniques:
Tampering, skimming, theft, pretexting, surveillance

Targeted victim

Industries:
Financial services, retail

Attributes:
Confidentiality, integrity, availability

Countermeasures:
CSC-17, CSC-20

Description
Peripheral tampering involves any tampering or physically manipulating a hardware device that connects to a computer system. These devices may include Personal Identification Number (PIN) Entry Devices (PEDs), scanners, printers, etc. Threat actors associated with altered payment card transaction devices typically involve organized crime groups.

* + tampering OR Skimmer (for 'All' only); no modifiers for 'payment card skimmers'

"Performance enhanced" PEDs.

"One of these things is not like the others, one of these things just doesn't belong ..."—Sesame Street

Detection and validation
This investigation involved the alteration of PEDs in a chain of stores. A PED is the hardware device one commonly sees sitting on the checkout counter of a merchant. One swipes the mag-stripe of the payment card, and if a PIN is used, one enters the PIN via the device keypad.

Response and investigation
On one particular occasion, the RISK Team responded to a request from a chain of retail stores involving possible altered PED devices. Our plan of attack was to travel to several stores collecting random samples of PEDs, which we then transported back to the RISK Labs.

Upon receipt, the Labs opened and visually inspected the PEDs. Their initial inspection of the devices immediately indicated that many had been altered. The devices in question had an additional mag-stripe card reader installed under the legitimate card reader. This allowed the mag-stripe of any card used to be read at the same time as the transaction. Additionally, a membrane touch keypad was installed under the legitimate device touch keypad. This allowed the PIN to be captured by the finger touches on the legitimate keypad during the transaction. An unauthorized circuit board was also installed in the altered PEDs.

Inspection of the extra circuit boards found a memory chip and a Bluetooth device. With some experimentation involving several of the altered devices, we were able to devise a process by which the memory chips in the skimming devices could be read. Once we determined how to read the data on the chips, we were off to the races.

Additional devices were collected from store locations identified through fraud reporting and transported to the Labs for examination. In a very short time, desks, conference tables and other work surfaces in the Labs were covered with dissected PEDs. In one day, we examined over five dozen PEDs.

The examination of the altered PEDs gave us a good idea of how they functioned. An altered PED was installed at a store checkout location. All transactions—which included mag-stripe data and PINs—were captured and stored in the memory chip. The Bluetooth device installed in the PED had a range of about 50 yards, thus affording threat actors the ability to connect to the device from a car in the parking lot via a Bluetooth connection. Subsequently, the threat actors could, after data download, clear the memory card and make room for new transaction data.

In a very short time, desks, conference tables and other work surfaces in the Labs were covered with dissected PEDs.

Remediation and recovery
Based on our findings, a number of actions were undertaken in order to address this situation. The store managers inspected all PEDs at their locations and removed all altered PEDs. Anti-tamper and tamper-evident modifications were made to all PEDs at all the merchant locations. Through the examination of the CCTV recordings, several of the organized crime group threat actors were identified and arrested.

How we got here
An examination of the security department's CCTV recordings found there were images of several of the PEDs being swapped by the threat actors. A threat actor first casually walked down the line of checkout lanes looking at each PED to determine which PEDs hadn't yet been altered. He then walked up to and leaned against the counter near the PED appearing to look at merchandise around the checkout lane. As he did this, he used one hand to disconnect and remove the PED and used his second hand to reach in his coat, pull out another PED, and install it on the PED mount. As the altered PED was installed, the threat actor hid the original PED under his coat.

The entire swapping process on average took less than ten seconds. The process and the actions of the threat actors at normal speed were almost undetectable.

Attack-Defend Card

Scenario 8 [Lethal]. Hacktivist attack—the Dark Shadow.

Data breach scenario

Frequency:* 3%

Sophistication level: 1 2

Composition:
Activist group

Incident pattern

Pattern:
Web app attacks

Time to discovery: D W M

Time to containment: H D

Threat actor

Motive:
Ideology

Disposition:
Unknown, Syria

Tactics and techniques:
SQL injection, phishing, backdoor, C2

Targeted victim

Industries:
Information, public, financial services

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-3, CSC-4, CSC-9, CSC-11, CSC-12, CSC-13, CSC-18

Description
Hacktivist (activist using hacking) attacks involve threat actors motivated by ideology in an effort to maximize disruption and embarrassment to their specifically targeted victims, For this reason, we consider hacktivist attack as a lethal data breach scenario.

* + actor.external.activists

Smoke on the Water [plant].

"I'll never drink water; that is the stuff that rusts pipes."
—W.C. Fields

Operational technology vs. information technology
Industrial Control System (ICS) is a term used collectively to describe various types of automated systems that control industrial production. SCADA, or Supervisory Control and Data Acquisition, is a type of ICS. In terms of systems, these are broken down into two categories: Information Technology (IT) —those systems required for business purposes, and Operational Technology (OT)—those systems required for industrial automation purposes.

Detection and validation
The RISK Team was retained by a critical infrastructure customer to assess their networks for indications of a security breach. This customer, hereafter referred to as Kemuri Water Company (KWC), was responsible for supplying and metering water usage over a number of neighboring counties. From the onset, KWC was adamant that no evidence of unauthorized access had been uncovered and that our assessment was being conducted proactively—as part of an ongoing effort to maintain healthy operations of their systems and networks. In scope were all of KWC's IT systems, which supported end users and corporate functions, as well as OT systems, which were behind the distribution, control and metering of the regional water supply.

Behind the scenes, KWC was a likely candidate for a data breach. Its internet-facing perimeter showed several high-risk vulnerabilities often seen being exploited in the wild. The OT end of the water district relied heavily on antiquated computer systems running operating systems from ten-plus years ago. Even more concerning, many critical IT and OT functions ran on a single AS400 system. KWC referred to this AS400 system as its "SCADA platform." This system functioned as a router with direct connections into several networks, ran the water district's valve and flow control application that was responsible for manipulating hundreds of Programmable Logic Controllers (PLCs), housed customer PII and associated billing information, as well as KWC's financials.. Moreover, only a single employee was capable of administering it. If a data breach were to occur at KWC, this SCADA platform would be the first place to look.

The OT end of the water district relied heavily on antiquated computer systems running operating systems from ten-plus years ago.

Interviews with the KWC IT network team revealed concerns surrounding recent suspicious cyber activity. It became clear that KWC management was aware of potential unauthorized access into the OT systems of the water district. More specifically, an unexplained pattern of valve and duct movements had occurred over the previous 60 days. These movements consisted of manipulating the PLCs that managed the amount of chemicals used to treat the water to make it safe to drink, as well as affecting the water flow rate, causing disruptions with water distribution.

Suddenly, it seemed there was a lot more to the story. We smelled smoke. And where there's smoke well, you know the rest. Our "proactive assessment" seemed more and more to be a "reactive investigation."

Response and investigation
With interviews still underway, KWC granted us permission to cross-correlate KWC's internet traffic against our Verizon Cyber Intelligence Center's cyber intelligence sources to identify evidence of communications with known threat actors. This test showed positive results. First, adversary IP addresses from

three recent investigations were found connecting to KWC's internet payment application. Second, these IP addresses were originally encountered while investigating hacktivist attacks. Finally, the suspect connections corroborated with the payment application's web server logs, suggesting likely exploitation of internet-facing vulnerabilities detected earlier in our assessment.

At this point, we had identified likely evidence of a security breach with an avenue of intrusion and had compelling IOCs based on first-hand precedent. Next step: prove or disprove our preliminary findings.

The internet payment application enabled KWC's customers to conveniently access their accounts from a laptop, a desktop system or even a mobile device. However, a quick look showed some serious security flaws. Access to customer water usage, PII and payment data required only a username and password. No second authentication factor was needed. Next, we found a direct cable connection between the application and the AS400 system. Making matters worse, the AS400 system had open access to the internet and its internal IP address and administrative credentials were found on the payment application webserver in clear text within an initialization (.ini) file. In other words, we found a high probability that any unauthorized access on the payment application would also expose sensitive information housed on the AS400 system. We clearly uncovered a fire trail.

As our findings later showed, the threat actors in fact exploited an easily identified vulnerability in the payment application, leading to the compromise of customer PII and payment information. The total population of unique records exfiltrated from the AS400 system exceeded 2.5 million. No evidence of fraudulent activity on the stolen accounts could be confirmed. That was the good news. The bad news was customer information was unfortunately not the full extent of the breach. The typical semantic footprint of a hacktivist attack shows greater interest in denying and disrupting the victim's ability to conduct business than stealing information for financial gain. That was definitely the case here.

Figure 2

1. Customer access to payment system
2. External (Internet)
3. Internal (Corporate access)
4. AS400
5. PLC management
6. Finance access (PII access)
7. IT management

Our endpoint forensic analysis revealed a linkage with the recent pattern of unauthorized crossover. Using the same credentials found on the payment app webserver, the threat actors were able to interface with the water district's valve and flow control application, also running on the AS400 system. We also discovered four separate connections over a 60-day period, leading right up to our assessment.

During these connections, the threat actors modified application settings with little apparent knowledge of how the flow control system worked. In at least two instances, they managed to manipulate the system to alter the amount of chemicals that went into the water supply and thus handicap water treatment and production capabilities so that the recovery time to replenish water supplies increased. Fortunately, based on alert functionality, KWC was able to quickly identify and reverse the chemical and flow changes, largely minimizing the impact on customers. No clear motive for the attack was found.

Remediation and recovery

With a clear indication of what occurred, working with KWC we turned toward remediation. Immediately IT administrators shut down access to and from the account management web front end and blocked outbound connectivity from the AS400 system. With the threat actor's access presumably cut off, systems were rebuilt with baseline images and placed back online. We recommended to KWC that it replace its older systems with modern versions and apply up-to-date patching as necessary. Additionally, we highlighted issues in its continuity planning revolving around single points of failure—namely the lone AS400 system administrator. In addition to having no backup for emergencies such as this, operating alone and without oversight, configuration choices made for convenient management were unchecked by security considerations.

Lessons learned

KWC's breach was serious and could have easily been more critical. If the threat actors had a little more time, and with a little more knowledge of the ICS/SCADA system, KWC and the local community could have suffered serious consequences. Large organizations are often complex and KWC was no different. On top of the complexities of actually managing and delivering water to homes, KWC also handled all of its own account and transaction information. In complex systems like these, centralizing assets can make management easier, but cannot be done without thought to security.

Having internet facing servers, especially web servers, directly connected to SCADA management systems is far from a best practice. Many issues like outdated systems and missing patches contributed to the data breach—the lack of isolation of critical assets, weak authentication mechanisms and unsafe practices of protecting passwords also enabled the threat actors to gain far more access than should have been possible. KWC's alert functionality played a key role in detecting the changed amounts of chemicals and the flow rates. Implementation of a layered defense-in-depth strategy could have detected the attack earlier, limiting its success or preventing it altogether.

Many issues like outdated systems and missing patches contributed to the data breach.

A walk down memory lane
ICS/SCADA systems have very specific purposes, such as controlling valves, regulating the flow of air and water and collecting data, such as temperature and flow-rates. At first, many ICS/SCADA systems were analog and air-gapped from corporate networks; the requirement for cybersecurity wasn't given any consideration.

As technology matured and ICS systems grew, there was an obvious benefit to move to more digital systems, centralized control rooms and networking to enhance efficiency. Cybersecurity was still not a consideration

As organizations grew, statistical information needed to be available in "real time," which required the once air-gapped ICS/SCADA network to be connected to the corporate or business network. Remote support became a significant cost savings and a convenience furthering the ICS/SCADA network connectivity directly to the internet. All of this connectivity has made ICS/SCADA systems vulnerable and much easier for threat actors to gain unauthorized access.

Many organizations acquired capital funds and purchased the latest and greatest, costing tens to hundreds of thousands of dollars, as a vulnerability mitigation strategy. Unfortunately, a strategic plan to move forward and maintain the new technology was often overlooked.

This new technology can provide a false sense of security, as operating budgets do not take into account the time to support, maintain and operate the new technology—thus it becomes ineffective. Threat actors have the upper hand when technology is not maintained and they develop ways to circumvent how it works. Continuous operational and security training, coupled with additional staff, are required to stay on the same level playing field as threat actors.

A defense-in-depth strategy is necessary to properly deter, prevent, detect and remediate. A good defense-in-depth plan will take advantage of people, processes, technology, awareness and technical training and a network environment that enforces isolation, logging and alerting.

Attack-Defend Card

Scenario 9. Rogue connection—the Imperfect Stranger.

Data breach scenario

Frequency:* 4%

Sophistication level: 1 — 2 — 3

Composition:
Organized crime

Incident pattern

Pattern:
Crimeware, privilege misuse

Time to discovery: W — M

Time to containment: M

Threat actor

Motive:
Financial

Disposition:
Varies

Tactics and techniques:
Phishing, C2, backdoor, spyware/keylogger

Targeted victim

Industries:
N/A (opportunistic)

Attributes:
Confidentiality, integrity, availability

Countermeasures:
CSC-2, CSC-3, CSC-8, CSC-12

Description
Rogue network devices range from wireless access points and personal laptops to any unmanaged asset connected to the corporate network. Organizations manage this risk using different controls, such as Network Address Control (NAC), 801.x authentication, separate Bring Your Own Device (BYOD) networks, and scanning for non-sanctioned or unknown devices. These unmanaged devices represent a significant threat as they can provide threat actors with an extremely flexible platform to enumerate and compromise the network in ways, which IT-managed devices may restrict.

*+ action.misuse.Use of unapproved hardware

BYOD'oh!

"Nobody can be as agreeable as an uninvited guest."—Kim Hubbard

Detection and validation
A customer in the finance industry contacted the RISK Team after receiving several complaints from their customers who weren't able to access their accounts through the customer website. When these account holders attempted to access their accounts, they received odd error messages indicating the site was blocked due to security concerns. Our customer, the victim, was unaware of any ongoing incidents and was in dire need of technical assistance to determine exactly what was happening.

A digital forensics/incident response urban legend
The internet is replete with stories of servers online and still in use unbeknownst to the IT security team even after being walled in by contractors remodeling an office building, or inadvertently left behind after a company moves out of a location and forgets to move that lone server in a closet in a remote corner of the building. We haven't seen this happen, but it wouldn't be shocking to find out there is some truth to the legends.

One thing is certain: We do often encounter victims who are disinclined to allow us to inspect certain systems and servers because they are supposedly off-line and/or have no data. Once all other avenues of infection are exhausted they reluctantly allow us to look at these assets, many of which are indeed still internet-facing, haven't been patched in years, and in some cases, are the initial means of the attack.

Response and investigation
We were met by the victim's IT security team lead, hereafter referred to as "Jill" after we arrived onsite. Jill quickly brought us up to speed on the response efforts. The IT security team had validated that the servers were operating normally and all anti-virus scans were coming back clean. On the surface, there were no obvious reasons for the issue at-hand, so we hunkered down and expanded the search to try to gather more details.

We began by querying various intelligence sources, both internal and open-source, to determine if the victim's IP address space was listed as malicious and more importantly—why. The results of this research were limited, but the IP address ranges of the site had been associated with C2 servers. This information was critical as it narrowed the search and revealed why the victim was having its IP address space blocked. We were able to provide Jill with potential IOCs to search within her network environment for signs of the malware and/or associated network traffic.

The initial review of the production network yielded no evidence of compromise. No traffic to known C2 servers was uncovered and even with updated signatures, no existence of malware was discovered. The organization had segmented guest and BYOD networks and after coming up empty on the corporate local area network (LAN), we worked with the network staff to review activity on the guest and BYOD networks—where we eventually hit pay dirt.

Looking through the network logs for the BYOD network, we found traffic to known C2 servers. Using these same network logs, Jill and the IT network team were able to identify the specific device—an employee's personal laptop. This system had been infected with malware at home, brought to the office and connected to

the BYOD network. So far, so good. The BYOD network exists to limit malware propagation from untrusted devices into the corporate network, and in this regard, as far as we knew, it succeeded as the corporate network was unaffected.

The problem arose in the mentality applied to these guest and BYOD networks. Since these weren't production or corporate networks, and no sensitive information was expected to exist on them, the company had implemented minimal controls and monitoring. This first oversight allowed the egress traffic to leave the network without any filtering – good for end users wanting to reach Facebook but bad for preventing malicious activity.

Figure 3

1. Internet
2. Firewall (NAT/Shared external IP addresses)
3. BYOD DMZ
4. Public WiFi
5. Corporate segment

Our review of the "whiteboard" network diagram also found that the guest and BYOD networks were going out through the same network equipment, as well as using the same Network Address Translation (NAT), as the corporate traffic. This resulted in the corporate network's reputation being affected by any-and-all devices connected to the guest and BYOD networks.

Remediation and recovery

Jill re-imaged the laptop and she also asked us for advice in preventing similar issues in the future. We pointed out that the BYOD and guest networks had worked flawlessly protecting corporate systems, but by sharing public IP address space unknown devices were able to influence the reputation of the company. Adjusting the network configuration to send this traffic out through a different interface would quickly resolve this issue. Jill wasn't satisfied with only mitigating this particular issue, however, and she quickly worked to add additional security around the guest and BYOD networks.

These changes included blocking high-risk network ports and protocols, and then implementing additional logging and monitoring. By adding these additional controls to her already segmented guest and BYOD networks, Jill was able to add even more assurance that these networks were safe for employees to use without risking the company's production assets.

Attack-Defend Card

Scenario 10. Logic switch—the Soup Sammich.

Data breach scenario

Frequency:* 53%

Sophistication level:

1 — 2 — 3 — 4 — 5

Composition:
Organized crime, unaffiliated, state-affiliated, activist group

Incident pattern

Pattern:
Web app attacks

Time to discovery:

○ — D — W — M — ○

Time to containment:

H — D — W — M — ○

Threat actor

Motive:
Financial, espionage, ideology

Disposition:
The U.S., China

Tactics and techniques:
SQL injection, use of stolen credentials, backdoor/C2, privilege abuse

Targeted victim

Industries:
Financial services, information, healthcare, public, education, retail

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-3, CSC-12, CSC-18

Description
Manipulation of account balances and withdrawal limits is a tactic used by financially-motivated threat actors to create non-existent funds, bypass security measures, and cash out quickly. These threat actors inflate banking and/or gift card balances by compromising databases, modifying specific tables (pump), and then initiating fraudulent transactions before losses are discovered (dump).

* + Hacking AND asset (web application)

"Pump and dump."

"Hello! We're back! I am Hans. And I am Franz. And we just want to ... [clap] pump ... you up!"—Hans and Franz, Saturday Night Live

Detection and validation
A Verizon Rapid Response Retainer service trigger came in from a customer in the banking industry. This victim reached out to the RISK Team after several of its high-value bank accounts had experienced millions of dollars of fraudulent ATM transactions. These transactions occurred in several countries around the world and involved massive withdrawal amounts greater than the standard allowable limits.

Response and investigation
The onsite investigation kicked off with bank security stakeholder interviews. We learned that the incident involved several hundred fraudulent ATM withdrawals at multiple locations globally—all within a two-hour window. These withdrawals impacted multiple bank accounts. The bank advised us that all of the accounts used in the apparent "pump and dump" transactions had a specific Issuer Identification Number (IIN).[5]

Our point of contact, hereafter referred to as "Bill," also passed along information that the bank had learned while working on the incident with law enforcement: one of the bank's IT administrators was involved. This administrator stole credentials from another administrator and then modified the security around the high-value accounts.

Using this information to focus our efforts, we collected system images and network logs to piece together the full story. The logs revealed that the malicious administrator had modified certain controls used to protect authentication information. Furthermore, the administrator fraudulently transferred money into the accounts and removed the withdrawal limits. Shortly after these changes were in place, the logs showed a large number of transaction requests being sent through the processing system over the course of minutes.

After providing Bill with an update, we looked for, and found, a piece of malware downstream from the device used to encrypt the transaction data. This malware captured a large number of encrypted transaction details and, because of the degraded encryption, was able to build a look-up table converting encrypted data to plain text. The same piece of malware had the ability to collect other PCI sensitive data. Combining this information allowed seemingly legitimate fake cards to be created and then used to liquidate these assets.

Once the key pieces of information were acquired, the threat actor struck and struck quickly. In two hours, the credentials were used to withdraw millions of dollars across numerous locations.

This malware captured a large number of encrypted transaction details and, because of the degraded encryption, was able to build a look-up table converting encrypted data to plain text.

5 en.wikipedia.org/wiki/Bank_card_number#Issuer_identification_number_.28IIN.29

Leveraging institutional knowledge
In many investigations, we find that someone within the victim organization may have a good idea of what might have happened, including what vulnerabilities may have been exploited. Often, these employees have previously suggested containment or remediation measures for observed vulnerabilities, but their input hasn't been taken into consideration. This can create frustrated employees who feel that, with their knowledge and experience, they could have helped prevent a situation if only someone had listened.

We always look for that person as the individual who has information can save time and help focus the investigation early on. In this case, an IT administrator advised us, based on his knowledge of the network environment, how this specific incident could have been accomplished. We were able to almost immediately focus on the key systems and quickly identify an administrator account that had been used to exploit the high-value accounts.

Remediation and recovery
Just like that, the deed was done. While Bill and his bank wrestled with the challenges associated with the stolen money, we worked with law enforcement and provided them with the additional details. Law enforcement later concluded that organized crime gangs had recruited IT administrators at multiple banks to install malware and to force configuration changes. Withdrawals were seen in several foreign jurisdictions in conjunction with these attacks.

We recommended to Bill that he identify roles within the bank that required access to sensitive customer data and implement monitoring controls around the data and the systems containing the data. In addition, we helped him create a review process to detect changes to encryption schemes and withdrawal limits. Bill also indicated the bank had plans to implement account balance increase monitoring that could detect unusual numbers of accounts being "pumped up." Using tactics like this, on top of established operational security practices, Bill and his bank will stand a much better chance of mitigating these types of advanced threats.

Leveraging insiders was obviously essential in this case, but our experience has shown that they aren't always necessary in "pump and dump" operations. Structured Query Language (SQL) injection attacks have been used to enumerate database schemas and manipulate the relevant fields to inflate gift card balances and compromise account numbers.

Configuration exploitation

Weak configurations occur from both a device and a network standpoint. From a device standpoint, these are the vectors of compromise. From a network standpoint, these allow for easy lateral movement after an established foothold, or as an avenue for data exfiltration.

At the risk of appearing to beat a dead horse, we would be remiss if we didn't bring up the fact that the top 2 hacking varieties take advantage of a weakness in static authentication mechanisms, specifically use of stolen credentials and password guessing (aka brute force). In fact, 80% of data breaches involve exploitation of stolen, weak, default or easily guessable passwords.

Web apps are the vector of attack in two of the scenarios in this section; the top industries with confirmed data breaches within the web app incident classification pattern are financial services, retail, information, utilities and public.

Attack-Defend Card

Scenario 11. SQL injection—The Snake Bite.

Data breach scenario

Frequency:* (11% overall) 23%

Sophistication level: 3

Composition:
Activist, organized crime, state-affiliated

Incident pattern

Pattern:
Web app attacks

Time to discovery: D W M

Time to containment: D W M

Threat actor

Motive:
Financial, espionage, ideology

Disposition:
Varies

Tactics and techniques:
SQL injection, backdoor, password dumper, use of stolen credentials

Targeted victim

Industries:
Utilities, manufacturing, public, education, retail, financial services

Attributes:
Confidentiality, integrity

Counter-Measures:
CSC-3, CSC-12, CSC-18

Description
SQL injection attacks, in their most basic form, are methods of abusing an application's interaction with its back-end database. These attacks leverage non-validated inputs to modify existing database queries to achieve unintended results and frequently target web applications.

* + SQL injection

A payday play day.

"Never spend your money before you have earned it."
—Thomas Jefferson

Detection and validation
A US-based industrial parts manufacturer contacted the RISK Team about a "possible" data breach situation. They qualified it as "possible" because they weren't really sure if a crime had taken place or not. They advised us that they had spent the previous two weeks looking for evidence of a breach, but were unable to find anything.

A cyber investigation firm on retainer: part of a sound escalation plan
Over the years, our customers have engaged us for various cyber-related incidents. These have included assisting their security and response teams, operating under parameters of attorney-client privilege and conducting PCI investigations, among other reasons.

A cyber investigative response firm on retainer can demonstrate to a board of directors, executive management, cyber insurance carriers and other interested parties, that the company is taking a proactive stance in its IR capabilities. The benefits of having a cyber investigative response firm on retainer include:

- Having a contract already in-place, eliminating the need to shop around for assistance while simultaneously responding to a cybersecurity incident.
- Quick response service-level agreements (SLAs) for phone support, investigator onsite deployment and emergency malware analysis.
- Leveraging proactive IR services, such as first responder training and mock incident table-top exercises.
- Access to cyber intelligence, expert advice and assistance, as well as other security services.

It appeared that for the last two bi-weekly payroll cycles, not one member of the company's C-suite had received their direct deposit paycheck. Instead, the paychecks were routed to a foreign bank account. In total, a number of the most senior management were directly impacted.

When it first occurred, the company ordered an internal investigation into the matter. The IT security team searched HR direct deposit databases looking for any evidence that data had been altered. It interviewed key payroll personnel and they examined perimeter firewall logs looking for something.

By the end of that first week, the IT security team found no evidence of a breach, concluding that the misrouted paychecks must have been the result of an error or glitch at the company's payroll bank. However, when the same thing happened two weeks later upper management wanted answers (and their money!).

Response and investigation
Given the internal politics and disagreement within the IT security team as to whether or not a breach had occurred, a critical first step was for us to identify a starting point for analysis. We asked the most basic questions—how does an employee interact with their payroll information? How do they update their direct deposit? We needed to understand what systems employees interacted with and

then determine if a threat actor was able to compromise those systems.

This is what we learned: This manufacturer maintained a web-facing HR portal that employees could access using their Social Security Number and a six-digit PIN as an authentication check. It was an internet-facing application with just about the worst authentication schema possible.

> Investigators: "When was the application built?"
>
> Victim: "Ten years ago."
>
> Investigators: "Where is the person who built it?"
>
> Victim: "Oh, he quit six or seven years ago."
>
> Investigators: "Okay then, when was the last time it was updated?"
>
> Victim: "Updated?... Hmm, it's never been updated."
>
> Investigators: "When was the last time it was subject to a vulnerability scan?"
>
> Victim: "Never."

At this point, we knew we were likely dealing with a SQL injection attack.

At this point, we knew we were likely dealing with a SQL injection attack.

A very basic external vulnerability scan indicated that certain elements of the HR website were susceptible to compromise via SQL injection. Along with basic login functionality, the non-authenticated side of the HR portal featured a "help" form where employees could create and submit basic questions. This was the form that was leveraged to interact with the database.

In reviewing the web server logs, we were able to confirm that it wasn't only attacked, but also successfully exploited tens of thousands of times!

The initial attacks were aimed at understanding the underlying database schema, structure and operational capabilities. The threat actors then moved on to manipulate the database to execute system-level commands in order to download additional tools from the web. With the database running with admin privileges, the threat actors were able to alter the direct deposit routing numbers, account numbers and bank information for members of the executive suite, whose information they had gleaned while enumerating database contents.

The victim stated, "We checked the HR direct deposit database. Everything was just the way it was supposed to be. We didn't find any altered records." This was true, the direct deposit information in the database was correct at the time that the IT security team examined it. This was due to the fact that after payroll was cut, the threat actors would go in and change the affected employees' checks back to their original values. And without actually examining web server logs, administrators would be none the wiser.

Remediation and recovery

As one might expect, once the organization had fully realized what had occurred, it opted to eliminate the HR platform and build a new one from the ground up – this time internally facing only.

Lessons learned
A key component that makes SQL injection interesting is that the vulnerability is actually based on poor web application implementation on the part of the victim, rather than vulnerabilities in the underlying database itself. The web server and database are essentially pawns in the attack and are just executing valid instructions—albeit with malicious intent. Therefore, there are no "patches" or "software updates" to rid the world of SQL injection. This can truly only be resolved through improved secure development principles and enhanced incident detection capabilities.

Gettin' in with the right, wrong or any PIN
A few years ago, we experienced another case involving a victim's payroll.

A large US-based retailer triggered the Verizon Rapid Response Retainer service at the tail end of the holiday shopping season. It was dealing with an odd discrepancy involving seasonal in-store workers not receiving their paychecks, even though these were being sent out by snail mail on a weekly basis.

Because of the transient nature of the workforce, the retailer issued payroll in the form of pre-paid debit cards. It afforded the retailer the ability to quickly hire (and fire) workers without having to address typical HR payroll issues, such as direct deposit. At first, only piecemeal complaints came in from employees, but a month later, hundreds of employees weren't being paid.

We started by doing a walkthrough of the debit card issuing process, up to and including the separate mailers, which consisted of first the debit cards, and then the PINs. We learned that even though the PIN mailers were on official company letterhead, the retailer wasn't mailing them. Instead, it printed thousands of letters, including the employee name, address, and employee ID, and shipped boxes of these to the issuer, that would then print the PINs onto the letters, and ship them out. In short, we found no indication at all that the relevant PINs were stored anywhere on the retailer's systems. Simply put, we were completely stumped.

As a last resort, one of our more seasoned investigators suggested we get one of the payroll debit cards and use it. The retailer arranged for access to a real employee payroll debit card. We walked over to an ATM at a supermarket adjacent to the retailer's corporate offices and attempted to use the card. We keyed in the PIN "1234" and successfully withdrew a $20 bill. We tried again, this time keying "5678." Again, we pulled out another $20. Then, just to re-revalidate, we keyed in "4321" as a PIN, and sure enough, we pulled out another $20.

It turned out that the authorization servers at the issuing bank were configured only to check that a 4-digit PIN was entered at the time of ATM withdrawal, not whether it was the correct PIN. The retailer quickly forced the bank to reconfigure its authorization servers and the criminal activity stopped immediately.

Attack-Defend Card

Scenario 12. CMS compromise—the Roman Holiday.

Data breach scenario

Frequency: 46%

Sophistication level: 3

Composition:
Organized crime

Incident pattern

Pattern:
Web app attacks

Time to discovery: D W M

Time to containment: D W

Threat actor

Motive:
Financial, espionage, ideology

Disposition:
China, Malaysia, the U.S., Russian Federation

Tactics and Techniques:
Backdoor/C2, use of stolen credentials, export data, SQL injection

Targeted victim

Industries:
Financial services, public, retail

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-3, CSC-12, CSC-18

Description
In today's network environment, content management systems (CMS) are ubiquitous. CMS provide an incredibly wide array of functionality including publishing, modifying content, organizing data and managing users. As with everything else, the more common it is, the more often it is targeted. CMS vulnerabilities that are left unpatched are often targeted and used as a foothold to install backdoor programs. A backdoor can lay in wait for days, months or even years before threat actors return to use it to gain access again.

Pirates on the high-seas.

"Knowledge is the treasure of a wise man."—William Penn

Detection and validation
The RISK Team was contacted by a global shipping conglomerate that advised they were having problems with piracy. Not software piracy. Actual piracy, as in criminals with boats and guns (although sadly, there were no parrots involved). The victim went on to describe how over the last several months, pirates had been attacking their ships traveling in shipping routes while on the high-seas. Obviously, piracy wasn't a new problem for this (or any other) shipping company. However, in recent months, the pirates had changed their tactics somewhat, and in a manner that the victim found extremely disconcerting.

Rather than spending days holding boats and their crew hostage while they rummaged through the cargo, these pirates began to attack shipping vessels in an extremely targeted and timely fashion. Specifically, they would board a shipping vessel, force the crew into one area and within a short amount of time they would depart. When crews eventually left their safe rooms hours later, it was to find that the pirates had headed straight for certain cargo containers. It became apparent to the shipping company that the pirates had specific knowledge of the contents of each of the shipping crates being moved. They'd board a vessel, locate by bar code specific sought-after crates containing valuables, steal the contents of that crate—and that crate only—and then depart the vessel without further incident. Fast, clean and easy.

Yet the question remained: How on Earth did pirates know which vessels to attack and which crates held their potential plunder?

Response and investigation
With this background information in hand, we began to enumerate where this type of information resided within the shipping company's systems environment. What we learned was that the company used a homegrown CMS to manage shipping inventories and specifically the various bills of lading associated with each of their shipping vessels.

We then honed in on the network traffic surrounding the CMS managing shipping routes. We discovered that a malicious web shell had been uploaded onto the server. The threat actors used an insecure upload script to upload the web shell and then directly call it as this directory was web accessible and had execute permissions set on it—no Local File Inclusion (LFI) or Remote File Inclusion (RFI) required. Essentially, this allowed the threat actors to interact with the webserver and perform actions such as uploading and downloading data, as well as running various commands. It allowed the threat actors to pull down bills of lading for future shipments and identify sought-after crates and the vessels scheduled to carry them. Fortunately, these threat actors made several mistakes, which we were able to capitalize on.

The threat actors used an insecure upload script to upload the web shell and then directly call it as this directory was web accessible and had execute permissions set on it—no Local File Inclusion (LFI) or Remote File Inclusion (RFI) required.

Common CMS attack chain of events
As you may have already guessed, most breaches leveraging CMS vulnerabilities don't involve maritime events. An increasingly common attack in our breach data set involves e-commerce applications being targeted and leveraging CMS weaknesses as the initial vector of attack. A typical CMS data breach involves an external threat actor performing the following steps:

1. Discovers an unpatched CMS application and leverages an RFI vulnerability to upload a web shell on the server, providing a backdoor and access to the web server.
2. Leverages the foothold to access the web application server.
3. Escalates privilege (if necessary) to fully compromise the web server and payment application.
4. Modifies the payment application code to collect payment card information as it is processed and dumps it into a file.
5. Leverages the web shell to extract the payment card dumps.

PCI requirements have helped drive secure payment card application coding that, for the most part, made clear text storage of payment card data a thing of the past. Criminals have changed their methods. Instead of hunting for PCI data, they're leveraging un-patched CMS applications and manipulating PCI-compliant application code to be non-compliant. This allows them to capture sensitive data, storing it for future collection or, in some cases, exfiltrating it in real time.

One of the first mistakes made by the threat actors was failing to enable SSL on the web shell. As such, all the commands were sent over the internet in plain text. This allowed us to write code to extract these commands from the full packet capture (FPC) data. We were ultimately able to recover every command the threat actors issued, which painted a very clear picture. These threat actors, while given points for creativity, were clearly not highly skilled. For instance, we found numerous mistyped commands and observed that the threat actors constantly struggled to interact with the compromised servers.

Our review of the FPC data revealed the threat actors tried, albeit in vain, to establish a reverse shell to directly interact with one of the compromised hosts. Try as they might, the threat actors were unable to move laterally. This attempt was blocked by a network security appliance. The threat actors then attempted to pivot to other systems within the network. They spent considerable time attempting to do so and, although armed with freshly dumped passwords, were unable to succeed. The threat actors also showed a lack of concern for their own operational security by failing to use a proxy and connecting directly from their home system.

Investigation spotlight: full packet captures
Full packet capture, or FPC, refers to using either network hardware or software to collect traffic for a given network segment. This traffic is captured in full fidelity with customer legal review and consent. Barring cases, such as HTTPS content encryption, it can give investigators unrivaled insight into an incident. This high level of visibility helps remove much of the guesswork involved during IR activities.

FPC systems can be used to extract very specific information from network traffic. Protocols such as HTTP or FTP use plaintext and easily parsable commands, allowing analysts to build exact timelines matched with specific actions. The fidelity of the collected packets also allows analysts—during an incident—to review prior traffic with new indicators.

Remediation and recovery
With all the information gathered, we were able to provide the victim with a clear and concise timeline of actions, compromised web hosts and data that was at risk. The victim was in turn then able to shut down the compromised servers, which—although important—weren't immediately critical to business operations. After blocking the threat actors' IP address, the victim reset all the compromised passwords and rebuilt the affected servers with current versions of its CMS.

Moving forward, the victim worked to adjust its security posture by starting regular vulnerability scans of its web applications and implementing a more formal patch management process. This would help mitigate the possibility of known vulnerabilities contributing to another incident. While these actions wouldn't prevent all attacks, they were certainly a step in the right direction.

Attack-Defend Card

Scenario 13. Backdoor access—the Alley Cat.

Data breach scenario

Frequency:* 51%

Sophistication level: 3 4 5

Composition:
State-affiliated, organized crime

Incident pattern

Pattern:
All

Time to discovery: D W

Time to containment: W M

Threat actor

Motive:
Espionage, financial

Disposition:
Romania, China, Russian Federation

Tactics and Techniques:
Backdoor/C2, phishing, export data, downloader, spyware/keylogger

Targeted victim

Industries:
Accommodation, financial services, public, professional

Attributes:
Confidentiality, integrity

Counter-Measures:
CSC-2, CSC-3, CSC-8, CSC-12, CSC-14, CSC-17

Description
Backdoors, along with C2 functionalities, are one of the most common footholds into internal networks. Once threat actors have a foot in the backdoor, they begin their post-compromise activities. Through this access vector, threat actors can now drop additional malware to perform a myriad of tasks, including capturing keystrokes, that lead to compromised accounts, escalated privileges, and movement to other areas in the victim's network, as well as establishing exfiltration points and methods for sensitive data.

* + malware (C2) OR malware (Backdoor)

What's your vector, Victor?

"Still round the corner there may wait, a new road or a secret gate."—J.R.R. Tolkien

Detection and validation
While conducting a routine review of network perimeter and other application logs for recent outbound connections, a customer IT security team discovered a problem. This customer, a manufacturing firm, found numerous instances of connections between the company's R&D department and an external IP address. This included numerous connections over the previous 24 hours involving an outbound transfer of over 2GB of data. Nothing on the network should have been transferring that much data! Alarm bells went off.

The victim raised its incident severity level from medium to high and, in accordance with its IR plan, triggered its third party cyber investigative response firm: The RISK Team. Cue Wagner's "Ride of the Valkyries"...

"Ask the data"
A peek into the incident data that feeds into the DBIR shows that over the last three years, for confirmed data breaches, "backdoor or C2" is the second most common hacking vector overall. During this same period, the malware functionalities of "backdoor" and "C2" are also both in the top five overall for data breaches.

When looking at data breaches involving an external threat actor an "espionage" motive, backdoor and C2 are over three times more likely to be associated with these data breaches than the breach database as a whole. The data shows a strong association between espionage-related data breaches and the installation and use of backdoor and C2 functionalities.

Response and investigation
Our investigation revealed conclusive evidence of a breach of an engineering team's shared computer system within the R&D department. As a result, user credentials for everyone who had used that system were compromised. With the customer's help, we were able to determine that the source of the breach involved a phishing email that targeted a specific individual on the engineering team. The phishing email resulted in a Remote Access Trojan (RAT) backdoor being downloaded onto the system, which enabled the threat actors to escalate privileges and capture user credentials.

Each of these steps positioned the threat actors closer and closer to their ultimate objective: To remotely access and exfiltrate a significant amount of highly confidential and proprietary information. This data represented months of R&D work and millions of dollars of investment by the manufacturer.

Leveraging the Verizon Cyber Intelligence Center and its IOC database, the foreign IP address and the malware were confirmed to be associated with a previously identified Advanced Persistent Threat (APT) group operating in Asia. This group was known to have perpetrated similar attacks against other US manufacturers. And so, the incident rolled out as follows...

Figure 4

1. Threat actor sends an email containing a link to a product site to an engineer's personal email address.
2. The engineer clicks on a link, which downloads and installs malware – a Poison Ivy variant – onto the engineering team's system. The malware installs a back door enabling threat actor access to the system, as well as capturing the keystrokes of system users.
3. Using the engineer's account credentials, threat actor accesses file shares within the R&D department.
4. Threat actor accesses the R&D department's accounts on cloud-based collaboration sites. Threat actor retrieves sensitive data stored on the file shares and collaboration sites.
5. Threat actor leverages a system as an exfiltration point for over 2GB of data over an FTP connection to an IP address in Asia.

Remediation and recovery

The damage was already done by the time this incident came to the victim's attention. The investigation revealed a significant amount of information had been leaked out via FTP to the foreign IP address. Remediation and recovery was relatively straightforward. However, as is often the case, it was achieved through a series of actions and on a rolling basis as details emerged during the investigation. As we do for many of the more complex and involved incidents, we broke down our recommended remediation and recovery measures into immediate-term and longer-term solutions.

Lessons learned

We must be realists and understand that 100% prevention of backdoor installation is as feasible as a "perpetual motion machine." That being said, the ease of post-compromise lateral movement was most concerning. The lack of common security controls for accessing and handling sensitive proprietary information in the R&D department segment and use of external collaboration services in an insecure and unauthorized manner represented serious issues in data control. The ability to establish FTP connections from a system regularly storing and processing sensitive data indicated insufficient monitoring. Finally, the overall lack of security awareness, including details surrounding sensitive projects, and especially in public social media settings, meant employees weren't aware that their actions were degrading the security of the project.

Most organizations do a good enough job of limiting the number of devices exposed to the internet. Threat actors are no less motivated to compromise assets on the internal network and will reach out and touch computer systems in RFC-1918[6]-land by leveraging the allowed services that provide interaction with the outside world. Whether the vector is a visit to a compromised website or a malicious email attachment, what follows is frequently malware installation that opens up a gateway providing ongoing access and control of the internal device.

Unfortunately, the combination of these factors made it easy to advance from a compromised system to intellectual property theft and made for a hard lesson in computer security for this company.

6 Request for Comment 1918 (RFC 1918) "Address Allocation for Private Internets" sets the IPv4 standard for assigning private (internal) network IP addresses.(for example, IP 10.0.0.1, IP 192.168.1.1).

The road to recovery: containment, eradication and remediation
In this particular case, we recommended several containment, eradication and remediation solutions to help the victim on the road to recovery.

The first set of recommendations were the "immediate-term" solutions—those that were critical, yet could be accomplished with minimal resources or within a short time frame:

- Implement network-level blocks of the foreign IP addresses.
- Block all internet access for the in-scope network segment.
- Force-change passwords for all employees.
- Prohibit further access to external, cloud-based collaboration sites.
- Pull the infected system off the network for forensic analysis.
- Conduct network-wide scans for malware signatures.

The second set were the "longer-term" solutions—those that were critical, yet required additional resources and/or time to implement:

- Implement two-factor authentication for access to the in-scope network segment.
- Employ a robust DLP solution, starting with the in-scope network segment.
- Prohibit employee access to personal email and other non-work related websites.
- Increase security awareness and employee training.
- Obtain executive-level attention to overall security needs.

Attack-Defend Card

Scenario 14 [Lethal]. DNS tunneling—the Rabbit Hole.

Data breach scenario

Frequency: <1%

Sophistication level: 3

Composition:
State-affiliated, organized crime

Incident pattern

Pattern:
All

Time to discovery:
Unknown

Time to containment:
Unknown

Threat actor

Motive:
Financial, espionage

Disposition:
Varies

Tactics and Techniques:
Export data, capture stored data

Targeted victim

Industries:
Retail

Attributes:
Confidentiality, integrity

Counter-Measures:
CSC-6, CSC-8, CSC-12, CSC-13

Description

The Domain Name System (DNS) enables people and machines to communicate across the internet. DNS essentially translates human-friendly domain names into machine-friendly IP addresses, and vice-versa. DNS tunneling is characterized by establishing an unintended communication channel to a C2 server and/or to exfiltrate data. This type of tunneling is often conducted on networks with strict security controls, as the pervasiveness of DNS means it is often allowed on highly restrictive networks.

The DNS protocol was never intended for data transfer (let alone for nefarious C2 activities), and as such, it is often forgotten about from a network security/monitoring standpoint. Manipulating DNS is nothing new, and is certainly something we have seen in various engagements over the years; nonetheless, it remains a lucrative means for miscreants to siphon off sought-after data. Because of this, we consider DNS tunneling as a lethal data breach scenario.

A two-lane superhighway.

"Two roads diverged in a wood, and I—I took the one less traveled by, and that has made all the difference."—Robert Frost

Detection and validation
A RISK Team customer had just been informed of a massive breach and had no idea where to look. This customer had been experiencing intermittent security issues for the past year, but there was no definitive explanation for this. What it had seen was a variety of drive-by infections and non-targeted phishing attempts, which was nothing out of the ordinary for an organization of this size.

When we arrived at the customer's location, we were quite happy to find it had a good handle on the situation. It immediately provided us with all recent anti-virus logs and incident reports on recent phishing attempts. We were also given access to various accounts on security appliances and log sources.

With support from the RISK Labs, we worked through the security events and logs available to us and ruled out obvious signs of compromise. While a few small incidents had popped up over the last year, they all seemed to be isolated events. After discovering one of the security appliances wasn't properly logging domain names, we turned to the DNS server logs to find domains related to specific events. Before we had a chance to review the security events, strange entries in the DNS logs jumped out at us.

Exfiltration tactics
One of the biggest problems with identifying exfiltration is the variety of techniques that can be used. While in the case of this customer, the DNS protocol was abused, the threat actors have sundry other options at their disposal. Frequently, free file sharing or cloud hosting services are used as exfiltration addresses and data is transmitted via HTTP or HTTPS. Other times the exfiltration host is a compromised internet server, a C2 server or a bulletproof hosting operation.

In many cases, it can be very difficult to determine whether the traffic is legitimate user activity or unauthorized exfiltration, without deep packet inspection or interviews with end users. All protocols have the potential to be abused and used to hide outgoing data. While monitoring is a necessary piece of the defense, the varied nature of exfiltration emphasizes the need for data segregation and at-rest encryption for sensitive information.

Response and investigation
Within the DNS logs were thousands of malformed, seemingly random entries, all sourced from just three IP addresses—backup servers used for the user account databases. The backup servers were on a secured piece of the network and had only limited access over SSH from designated administrators and DNS resolution through internal name servers. The segmentation controls protecting the backup servers caused them to not be included as part of the original focal point of the organization. The upside was, after validating the firewall rules, we were able to narrow down the investigation to a very small subset of hosts.

Forensic images were collected from these now suspicious hosts and sent to the RISK Labs for analysis while we continued investigating the head-scratchingly-odd DNS traffic. We returned to the DNS logs, but rather than looking at the requests themselves, we looked at what the DNS server was doing with those requests.

All of the requests from the backup servers ended up being routed to a single remote name server. Each request included a domain rendered in hexadecimal characters of equal length and was met with the expected "invalid" response. To most network devices, this traffic would appear as "normal." An invalid request met by an "invalid" response was hardly a security red-alert, but the patterns showed strange things were afoot on User Datagram Protocol (UDP) port 53. Especially of concern was that this represented a direct path from the backup servers to the public internet.

An invalid request met by an "invalid" response was hardly a security red-alert, but the patterns showed strange things were afoot on User Datagram Protocol (UDP) port 53.

With a stick of bubble gum, a paper clip and a little bit of coding magic, we converted the invalid requests into a variety of other text formats. The results seemed entirely random until a text header of a ZIP file appeared. While possible, it was very unlikely this was purely random, so we modified the code to search for other ZIP file headers and results began to finally emerge. Was this DNS traffic actually data being exfiltrated? Log examination results also verified extractable compressed archive files containing the account backups could be retrieved from these DNS requests.

The RISK Labs reported results from the hosts with access to these backup servers that included a network administrator's desktop system with multiple pieces of malware installed. In the case of both the backup servers and the infected system, the malware had cleaned up local files and erased evidence necessary to build out a specific timeline. Regardless, we now had enough information to move forward with remediation.

Remediation and recovery

Our immediate response was to rebuild all affected systems and to implement new controls around the internal DNS servers. Backup servers were no longer allowed to make DNS requests for non-internal domains and security event systems were now being fed DNS logs to correlate. The anti-virus solution that had identified previous malware infections had missed this one due to disabled updates, an issue the customer resolved with both training and group policy changes.

In addition to the items above, we recommended that the customer:

- Periodically—especially after incidents—reviews and updates network Intrusion Prevention System (IPS) and firewall rules to prevent recurring issues.
- Continues to monitor for the previously-alerted activity, as well as any additional IOCs.
- Implements a File Integrity Monitoring (FIM) solution to monitor executable and other usual suspect file activities.
- Employs an application whitelisting solution that prevents unauthorized software from running on critical systems, to include those that store sensitive information.
- Implements a DLP solution to detect and prevent unauthorized attempts to copy or move sensitive data without authorization.

Lessons learned

The organization has a better understanding of how adversaries will exploit the limited services provided to them to reach back out to the internet, even if the services aren't designed for file transfer. Direct access from or to the internet was eliminated and mundane network traffic, such as name resolution, is now inspected and scrutinized to detect covert channels. While it is preferable to detect and stop an attack earlier in the kill chain, identification of new ZIP or RAR files and threat modeling paths of data escape can hinder exfiltration efforts and hopefully limit the breadth of the breach.

Malicious software

For us, it's almost cringeworthy hearing an incident described as "a malware incident." Over 50% of confirmed data breaches over the previous three years have featured malware somewhere along the attack chain. Whether malware is introduced as the very first event or used post-compromise to advance the attack (or both), it is expected to be an integral part of any sophisticated data breach. Data breaches involving malware also feature hacking actions 80% of the time and social actions 44% of the time. In other words, it’s typically not just a "malware incident."

"Ask the data"
In the real world, early detection can limit the impact of a malware infection and possibly break the event chain before data loss is realized. Below are the top functionalities or varieties of malware in our breach corpus, together with potential mitigation tactics:

1. Exported data – identify unsanctioned encryption and egress filtering/ monitoring.
2. Keyloggers – implement multi-factor authentication to limit effectiveness.
3. Backdoors – use proxies for internet-based traffic; implement packet inspection to detect use of port 80 traffic.
4. RAM scrapers – detect evidence of new data stores; identify new processes on critical devices. These are prevalent in POS compromises.
5. C2 platforms – use threat intelligence to blacklist known C2 servers.

Attack-Defend Card

Scenario 15 [Lethal]. Data ransomware—the Catch 22.

Data breach scenario

Frequency:* 4%

Sophistication level: 1 2

Composition:
Organized crime

Incident pattern

Pattern:
Crimeware

Time to discovery: H D

Time to containment: H D

Threat actor

Motive:
Financial

Disposition:
Varies

Tactics and techniques:
Ransomware, backdoor/C2, phishing

Targeted victim

Industries:
N/A (opportunistic)

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-8, CSC-10

Description
Ransomware is characterized by malware that prevents users, typically through encryption, from accessing their system, file shares or files. After gaining access and control, threat actors hold the data for "ransom" until the user agrees to pay money to regain access to their data. For this reason, we consider data ransomware as a lethal data breach scenario.

*+ malware (ransomware) in all incidents, not just data breaches

Data ran somewhere. Or did it?

"I, with my partie, did lie on our poste, as betwixt the devill and the deep blue sea."—Robert Monro

Detection and validation
In one such case a network administrator for a small company, hereafter referred to as "Sam," began receiving communications from users reporting issues with accessing a financial database application shared among multiple employees. Additionally, some users reported a strange webpage indicating their files had been encrypted and that they would need to follow specific instructions involving sending money for a decryption key.

Sam began going over his plan for getting the business back on track. The first action he took was to restore the systems to an image from backups. After wrestling with some issues, Sam discovered that due to hardware failures on the storage devices, the last known good backup was over three months old. Needless to say, this was unacceptable for production use.

Sam escalated the matter to his executives. After discussing the matter, the executives decided to acquiesce and pay the ransom to regain the keys to their kingdom. After converting cash into bitcoin currency in the amount demanded, they transferred the bitcoins to the wallet address presented on the system screen pop-ups. They then attempted to access the hyperlink provided, only to find that the web page had been taken down and there were no instructions on how to receive the keys. At this point, Sam, on behalf of his company, engaged the RISK Team. We got moving and quickly arrived onsite.

Response and investigation
The first order of business was to help Sam "stop the bleeding" and contain the situation. We were faced with ransomware on three systems, one of which contained the database application accessed by numerous users. After instructing Sam to leave the devices powered on, but disconnected from the network, we had Sam focus his attention on gathering additional data, to include network logs covering the in-scope time frame.

Our investigative approach included a typical plan of attack: automated malware scanning, timeline analysis, physical memory analysis log file examination and other activities. We soon identified that several thousand files across the systems had been encrypted by the ransomware. The systems indeed had internet access (which, according to the administrator, wasn't required) and the vector of compromise was a malicious Adobe Flash file that was accessed while online. It also appeared that as soon as the file (determined to be identical based on a hash value comparison) was downloaded to the systems, it replicated itself with a different file name in a different directory. File execution analysis confirmed that this file was executed just before the user files were encrypted.

Our investigative approach included a typical plan of attack: automated malware scanning, timeline analysis, physical memory analysis, and log file examination, and other activities.

We collected a listing of IP addresses sourced from network logs and provided them to the Verizon Cyber Intelligence Center for comparisons against their IOC database. This identified several dozen IP addresses associated with previous malicious activity. Furthermore, one of the IP addresses was affiliated previously with Adobe Flash related malware, a known transportation vehicle for this ransomware.

As it turned out, this specific version of the ransomware downloaded a common scripting language file prior to the payload being dropped. Moreover, the network activity we identified in close proximity to a text string found in the network logs, which included artifacts related to Adobe Flash files, reinforced the vector of compromise. Unfortunately, the settings for the default internet browser were set to clear its history, cookies and cache upon exit, so browser analysis failed to bear any fruit.

Malware spotlight: exploit kits
Exploit kits are a specific type of malware, consisting of pre-compiled exploits and are used to take advantage of vulnerabilities found in many heavily used software applications. These may include Adobe Flash, Java and Internet Explorer. Many exploit kits use rotating control domains and hourly anti-virus detection checks to avoid detection, thus defeating traditional blacklist and anti-virus solutions. These exploit kits often include control panels to make managing campaigns easier.

While the end goal—delivering a malicious payload—is the same, the process of how it achieves its goal can differ from one exploit kit to another. The problem today is that these toolkits are scalable, efficient, profitable and hard to detect. These toolkits can be relatively easy to use, appealing to a wide range of threat actors and tend to be responsible for the majority of infections worldwide.

At this point in the investigation, we had identified the malware and the vector of compromise; however, due to one common username and password being shared by all of the systems' users, we could not attribute the activity to any single user.

With data recovery hopes dimming, a breakthrough presented itself through binary and memory analysis. Where the more advanced and distributed ransomware networks kept their keys on remote servers, the threat actors in this instance used a rudimentary setup with the keys were stored in memory. We pulled the keys from the systems and provided them to Sam for running against the batch of files affected.

Remediation and recovery
We communicated our investigative process to Sam throughout and provided him with the latest IOCs, including malicious IP addresses, executable files and other relevant artifacts. Sam used these IOCs to scan his network in an effort to identify additional compromised systems (but this yielded negative results).

In terms of future prevention and mitigation, we provided Sam with several recommendations. These included removing internet access from critical systems (or, if required, changing internet browser preferences to track user history and cache); using individual user logins with a complex password policy; and ensuring applications, including Adobe Flash and JavaScript, are updated and regularly patched.

Digital forensic artifacts indicated that the ransomware was downloaded via a drive-by download that leveraged an unpatched Adobe Flash vulnerability. Our ability to perform forensic analysis on the in-scope systems and network logs, as well as tap into the Verizon Cyber Intelligence Center's extensive IOC database, provided the actionable intelligence necessary to identify the type, root cause and scope of the compromise. Fortunately, this compromise was limited to data on three systems, which, in the end, was recoverable.

Stop the bleeding, remove the shrapnel
Often times, due to business impact, servers may not be able to be rebuilt right away. If this is the case, typical remediation activities are as follows:

- Disconnect the network cable.
- Stop the malware process.
- Remove persistence mechanism(s).
- Delete the malware and any associated files.
- Restart the remediated system.
- Monitor for suspicious activity.

Lessons learned
With threat actor tactics evolving faster and faster, conventional security tools such as intrusion detection, log aggregation and anti-virus products are no longer enough. Having packet sniffers hunt for malware and other IOCs can make the difference.

Attack-Defend Card

Scenario 16 [Lethal]. Sophisticated malware—the Flea Flicker.

Data breach scenario

Frequency:* 32%

Sophistication level: 4 5

Composition:
State-affiliated, organized crime

Incident pattern

Pattern:
Cyber-espionage

Time to discovery: W M

Time to containment: W M

Threat actor

Motive:
Espionage

Disposition:
Varies

Tactics and techniques:
Backdoor, C2, export data, password dumper, spyware/keylogger, rootkit, exploit vulnerability, scan network

Targeted victim

Industries:
Varies

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-8, CSC-9, CSC-13

Description

Sophisticated malware attacks are best defined as those situations involving anti-virus and other security solutions that are rendered ineffective due to custom written, and often specially purposed malware deployed by threat actors. Sophisticated malware incidents typically take longer to detect, and tend to challenge the most mature organizations whose security controls effectively handle the simpler and well-known malware attacks.

As can be expected, sophisticated malware tends to occur more commonly with targeted attacks. While sophisticated malware is found in a smaller percentage of the breaches in our entire data set, it is present in a significant portion of our investigations. Because of this, we consider sophisticated malware as a lethal data breach scenario.

* + malware (3 or more varieties)

Premier cru maliciel.

"The more sophisticated we get, the more advanced our buildings and vehicles become, the more vulnerable we are."
—Stephen Ambrose

Detection and validation
A good illustration of the challenges posed by sophisticated malware attacks can be seen in a situation involving a multinational financial institution. The RISK Team was engaged to support the ongoing investigation. This financial institution believed that its intellectual property and customer information might have been compromised. The firm's executives had long since worried about insider threats and had focused their investigation entirely on behavioral analysis of key employees. Our involvement began after the firm's analysis yielded nothing indicative of insider misuse.

Response and investigation
Based on what we knew to date (and on the flip-side, an understanding of what we didn't know), we began by employing some broad-based network forensic data collection from the environment. We felt that the behavioral analysis that the firm performed was solid, but its lack of any conclusive results led us to believe that maybe the problem wasn't a "rogue" insider after all.

Our initial perimeter traffic capture produced some interesting findings in just the first few hours—traffic that was flagged as matching previously identified C2 servers on the internet. These C2 servers had been responsible for other intellectual property theft cases—so contextually, it was also interesting. This not only gave us the first breadcrumb to follow, but it also helped to narrow the investigation around a focal point within the firm's environment. Rather than trying to "boil the ocean" and attempt to whittle down hundreds of potential employees that might have had access to the data across dozens of offices around the globe—we had a single suspect IP address.

Unfortunately, the initial IP address was an external NAT[7] address and thus not tied uniquely to an individual system. The customer confirmed that it was indeed an IP address used for internet-related end user system traffic. We then hit an investigative speedbump in that egress traffic wasn't being logged; we had no means to correlate NAT'd traffic to an internal system. At this point, we had a solid lead, but just needed better instrumentation to improve the picture. Our next step was to deploy our custom-made network forensic appliances in order to begin collecting data and analyzing traffic patterns.

Upon set-up, these appliances soon lit up like Times Square on New Year's Eve; the results were astonishing. Not only did we quickly re-identify the suspicious egress traffic to the C2 servers, but we were also able to piece together a web of communications taking place internally among several dozen systems—many appearing to be compromised end user systems. The key anomaly was the unusual amount of port 53 traffic between the systems. The focus on the single IP address then expanded to these systems and forensic disk images were collected for analysis.

We were also able to piece together a web of communications taking place internally among over 40 systems—many appearing to be compromised end user systems.

The browser cache data on one of the systems helped us determine that the threat actors had gained initial access via social engineering. One of the personal emails on the system that was viewed contained a hyperlink associated with a malicious website. Upon clicking this link, installation of some basic system exploitation tools occurred, followed by lateral movement within the environment and installation of additional malware.

7 NAT (Network Address Translation)—used in this instance to translate the address of a group of internal devices to a single public IP address when connecting to the internet.

This additional malware set-up a listener on port 53 to impersonate a DNS daemon—the firm wasn't blocking (or logging) port 53 traffic anywhere in its environment. The malware operated in a peer-to-peer (P2P) fashion and exclusively from within RAM. Several of the compromised systems that were part of this P2P mesh were those of end users who were responsible for handling the sensitive data that had been the focus of the investigation.

Malware spotlight: peer-2-peer communication
Like many file-sharing clients, malware authors have learned to take advantage of the benefits of P2P principles to make malware more robust. P2P malware is able to avoid using a centralized control server for some or all of its functionality, allowing it to defeat traditional blacklisting. This decentralized "mesh" network is also highly resilient against hosts being identified and rebuilt, as still-infected systems can find a new infected endpoint with which to communicate.

The threat actors could leverage any of the compromised systems as a pivot point to expand the mesh if needed and/or re-infect systems if the malware was cleared out due to a reboot. In addition, the P2P capability provided an almost high-availability data exfiltration pathway, as any compromised node could be used to communicate out to any of the internet-based C2 servers.

In the end, through in-depth analysis we confirmed that after the threat actors gained access and conducted a thorough reconnaissance, they developed malware to extract data by targeting running processes related to the victim's operations (and, you guessed it, since this was customized malware, the resident anti-virus solution fell flat on its face).

Our malware reverse engineering determined that the malware contained various entropy settings to randomize its communications through encryption usage thus avoiding DLP. The data also went through a two-step validation process for authentication. Only after the data passed this validation processes was it written to an obfuscated output file. Talk about efficiency in coding!

Our malware reverse engineering determined that the malware contained various entropy settings to randomize its communications through encryption usage thus avoiding DLP.

Remediation and recovery
It was evident from the reverse engineering that this was no "garden variety" malware; a number of decisive and quick measures needed to be put into action to deal with this threat. The firm immediately enabled logging of all firewall activity. Ingress and egress traffic related to the internet C2 servers was blocked—cutting the malware off from the outside world and inhibiting its ability to reconfigure or spread. Lastly, internal Access Control Lists (ACLs) were implemented to drop port 53 traffic that was not destined for the firm's proper DNS infrastructure. As a result of this, network traffic spiked around compromised systems as the malware tried to re-establish connections.

We finally forced a reboot of the compromised system and watched the logs closely as they came back online. After a few days of monitoring, we found no unusual traffic and concluded that the threat had been removed.

Lessons learned
This victim faced sophisticated threat actors with a malware capability to match. The victim was so heavily focused on the possibility of a deliberate insider threat that it overlooked some of the foundational elements of incident detection within its security program. The aspects that it was examining were purely that of human behavioral anomalies, despite the fact that the unusual internal DNS traffic should have been visible had anybody been looking. Furthermore, the fact that the C2 destinations were known for months gave rise to additional conversation regarding improving the victim's threat intelligence posture as well as its processes for responding based on that intelligence.

In actuality, the lack of findings from the behavioral analysis was a "win" in the investigation, as it didn't allow confirmation bias to extend efforts into that specific threat scenario.

Attack-Defend Card

Scenario 17.
RAM scraping—the Leaky Boot.

Data breach scenario

Frequency:* (8% overall) 55%

Sophistication level: 2 3

Composition:
Organized crime, state-affiliated

Incident pattern

Pattern:
POS intrusions

Time to discovery: W M

Time to containment: D W

Threat actors

Motive:
Financial, espionage

Disposition:
Romania, Germany, China, Russian Federation

Tactics and techniques:
RAM scraping, brute force, export data, use of stolen credentials

Targeted victim

Industries:
Accommodation, retail

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-8, CSC-14

Description
RAM scraping involves malware designed to monitor and extract specific, targeted data from physical memory. Our analysis of this type of malware indicates that threat actors customize their tools to work in specific environments, such as on application-specific POS servers and terminals. Typical threat actors include those seeking financial gain, such as organized criminal networks, as well as independent hacking groups.

* + malware (RAM scraper)

RAM scraping. Ewe ...

"Do not trust your memory; it is a net full of holes; the most beautiful prizes slip through it."—Georges Duha

Digital evidence
From a digital forensics standpoint, collection of various evidence sources is necessary to maximize investigative findings. These evidence sources include current running processes and open network connections, also referred to as volatile data; a full, physical memory dump, which may contain information about recently run malicious processes; and a full sector-by-sector disk image, which may contain critical artifacts in active file space, as well as "deleted" files recoverable from unallocated disk space.

Detection and validation
When it comes to being notified externally of a PCI breach, many companies are typically contacted by one or more of three external entities: the payment card brands, their merchant bank or law enforcement. If the payment card brands provide notification, this is normally because the organization has been identified through fraud analysis.

This fraud analysis indicates multiple fraudulent transactions using cards previously used at a common location. In the case of law enforcement, investigators may be conducting an investigation into another matter and evidence comes to light that the retailer has been compromised. This case began with the victim being notified by a third party that it had been breached. The victim reached out to us, the RISK Team, a PCI Qualified Security Assessor (QSA) and PCI Forensics Investigations Firm, to investigate.

Response and investigation
Based on our substantial experience with POS intrusions, we focused on identifying any malware found on the payment processing systems – specifically RAM scrapers. While this type of malware has been around for several years, it can still be a challenge to detect and the information provided by the third party proved to be invaluable. Ultimately, the customer needed to know how the threat actors got in and how to prevent them from doing so again.

Using the intelligence provided by the third party along with various network logs provided by the customer, we quickly narrowed in on a number of systems that showed signs of malicious activity. Many of these logs came from perimeter protection devices alerting on, but not blocking, communications to suspicious remote addresses. The resulting list of suspicious servers quickly formed our marching orders for the evidence collection processes.

Using memory dumps and disk images collected from the systems, we were able to quickly locate a myriad of suspect processes. A memory analysis tool allowed us to extract the hidden processes and determine if they were the culprit. Within these processes we found a number of backdoors as well as a tool for dumping passwords, all left behind by the threat actors. These tools had allowed the threat actors to create administrator accounts and traverse the victim's network.

Based on our substantial experience with POS intrusions, we focused on identifying any malware found on the payment processing systems – specifically RAM scrapers.

The privileged user compromise allowed the intruders to gain access to processing servers, which handled all the victim's credit card transactions. These servers were then infected with a second, specialized piece of malware. Our analysis revealed this program as a RAM scraper with exfiltration capabilities. This particular piece of malware had been installed to hook into the specific process that took the credit card information as it was swiped.

With PCI requirements and common sense dictating that storing plain text payment card data is unacceptable, threat actors seek to capture this data in memory space, which is the only window of exposure before it is encrypted for transmission. After harvesting the physical memory of the infected device, credit card data was written to an encrypted file. This file would later be exfiltrated to a foreign IP address, confirmed in the network logging.

Threat targeting
Over the years we have seen many RAM scraper cases. From a threat-targeting perspective, simple strings searches against the associated malware files often leads to interesting right-off-the-bat finds, lending credence to the fact that the threat actor had done their homework. In one particular case, we pulled out text strings that revealed specific targeting across three areas:

1. The victim itself—yes, the actual victim name was embedded in malware.
2. The POS application—in the form of specific POS application file paths.
3. The payment card data—four payment card track data search expressions.

By combining all the log information along with the behavioral analysis of each piece of malware, we were able to establish that the customer had been initially breached two weeks prior through a targeted phishing mail. The threat actors gained remote access to the victim's servers using stolen credentials to access restricted devices. These devices were then used as beachheads for the threat actors to download additional tools to map the network. Once reconnaissance had been done, the threat actors targeted the processing servers they needed and installed customized RAM scraping malware to steal payment card information.

Remediation and recovery
This attack was complex, using a variety of tools and techniques to ultimately make off with a large stash of credit card information. The RAM scraping malware used on the processing server had components that were custom-built to extract the information desired and could prove to be very hard to detect. In this case, the IOCs provided by the third party saved valuable time in identifying where to look.

Having determined that there were sizable encrypted dumps that had yet to be transferred outside of the compromised device, the containment plan had to be quick and effective. The threat actors made multiple attempts to re-enter the victim’s network to retrieve said files. In the end, the customer had to completely disconnect from the internet while it rebuilt its network from the ground up over several days. This rebuild included properly segmenting its previously flat network, with two-factor authentication and strict firewall rules. In addition, the victim rebuilt the affected servers and reset every password in the environment.

"Ask the data"
POS intrusions have been dominated by one of two malware functionalities over the years – RAM scrapers and keyloggers. Both capture the desired data in process as the user is entering the information and before other security controls can be established.

Back in the 2012 DBIR days, keyloggers were the preferred tool, found in almost three quarters of breaches in the POS intrusion attack pattern (involving malware). Fast forward to the 2015 DBIR and RAM scrapers were identified in over 95% of POS intrusion breaches (involving malware).

Lessons learned
Many POS intrusions are what we call smash and grabs; the POS server is internet-facing and is the first door kicked in (many times with weak or default credentials). This particular incident was more sophisticated and targeted in nature. The POS environment wasn't the initial entry point; a blended attack method that included social (engineering), malware and hacking threat actions was utilized to gain access to the desired data. RAM scraping was front and center, and is a malware functionality found in most POS breaches large and small. Stolen credentials were leveraged in multiple stages of the attack and two-factor authentication, while not a panacea, forces threat actors to drastically adjust their tactics.

Attack-Defend Card

Scenario 18. Credential theft—the Poached Egg.

Data breach scenario

Frequency:* 42%

Sophistication level: 2 3 4 5

Composition:
Organized crime, state-affiliated

Incident pattern

Pattern:
All

Time to discovery: W M

Time to containment: H D W M

Threat actor

Motive:
Financial, espionage, ideology

Disposition:
Ukraine, China, Romania, Germany, Russian Federation, the U.S.

Tactics and techniques:
Use of stolen credentials, export data, spyware/keylogger, phishing, backdoor/C2, password dumper

Targeted victim

Industries:
Financial services, public, retail, professional services, information

Attributes:
Confidentiality, integrity

Countermeasures:
CSC-8, CSC-16

Description
Spyware/keylogger attacks involve unauthorized software or hardware introduced to a system to record user and system-generated information. Threat actor motivations range from financial, to espionage to ideology.

*+ hacking (use of stolen credentials)

Operating "inside the wire."

"I know imitation is the highest form of flattery, but stealing one's identity is totally different."—AnnaLynne McCord

Generally speaking, if a federal law enforcement entity contacts you it isn't because something awesome happened. This is an all too common method of breach discovery: law enforcement investigates one breach and identifies another victim of the same threat actors. To complicate things, law enforcement is often limited in what it can provide to potential victims. Naturally, this can send an organization into a frenzy trying to find the missing puzzle pieces.

Recipients of this not-awesome-at-all news soon experience the initial physiological responses—neuroendocrine activation, sweating profusely and heart pounding. Once the initial shock wears off, victims must consider the task at hand—an investigation involving in-depth forensic analysis to piece the puzzle together.

Detection and validation
The RISK Team was engaged by a victim as part of its IR plan escalation process to forensically investigate a potential compromise of its network. This victim had been contacted by law enforcement, who informed it that several of its external-facing IP addresses were involved in malicious activity associated with an active investigation. Representatives of the victim were provided with a list of suspected IP addresses and were left to determine what, if any, "malicious activity" was actually occurring on its network. We'll refer to our point of contact as "Ned."

Response and investigation
Upon arriving onsite, we met with Ned and soon began collecting forensic evidence from the suspected compromised systems, which included numerous web servers, domain controllers, and administrator and end user systems. Our forensic analysis soon identified that the systems were peppered with numerous variants of malware. We determined that the compromise spanned far beyond the systems originally identified by law enforcement. Our analysis of network log data showed that several other systems had been communicating with an external malicious IP address as well. We circled up with Ned, and gave him an update.

So, how could this happen in the first place? Our investigation revealed that the initial compromise occurred several months prior on one of the victim's webservers. Yes, you read that right; the threat actors went undetected and had free rein of the network for several months! We determined that the initial vector of compromise was a malicious web shell that was "dropped" onto the webserver.

Analysis of the web log data indicated that the threat actors had performed an SQL injection attack to gain access to the system via one of the victim's web-facing web servers. They then uploaded a web shell as a backdoor to continue the attack by executing arbitrary commands on the system. Once the box was compromised, the threat actors utilized malware of various functionalities and installed additional tools needed for the job.

In this attack, the compromised web server wasn't the end game. Malware capable of dumping passwords, stealing credentials via keylogging and performing exfiltration of data were all uploaded via the web shell. The adversary was aiming to penetrate deeper into the network using malware specifically designed to leverage existing credentials. Moreover, keylogging malware is made no less effective by the legacy focus on password length and "$peci@l" characters.

The compromised web server wasn't the end game. The adversary was aiming to penetrate deeper into the network using malware specifically designed to leverage existing credentials.

Figure 5

1. Foreign threat actor
2. ColdFusion web server
3. Admin password repository
4. AD domain controllers
5. Other network resources

The threat actors were successful in compromising several administrator accounts using password dumping, password cracking and keylogging. As our forensic analysis continued, we soon realized that the breadth of credential theft was much larger. We updated Ned.

In addition to log data and IP addresses, we also recovered a deleted spreadsheet that included a listing of more than 100 usernames and passwords for a specific administrator. As it turned out, this spreadsheet had been exported from a password repository service that the administrator was using to manage personal and professional passwords.

Two-factor authentication wasn't required to access the password repository service. This set of compromised credentials led to access gaining even more credentials and ultimately provided the threat actors with access to numerous network locations, including several domain controllers. Given the amount of time the adversary had rooted around in the environment, coupled with the effectiveness of the keylogging code, the amount of data that was at risk was significant.

Remediation and recovery

Given all of this information, there were a number of containment and remediation actions that Ned would need to take. First, he needed to attempt to eradicate any-and-all malware present on the victim's network. As such, we provided Ned and his team with specific IOCs that included malicious IP addresses, executable files, process names and other artifacts that they could use to scan their network.

In terms of remediation, we recommended that Ned force password resets across his entire company and emplace a password policy requiring a unique password per type of resource (an oldie but goodie). We also recommended that Ned implement two-factor authentication for remote VPN connections and, in particular, those connections accessing sensitive resources. Finally, we impressed upon him the need to improve his vulnerability patching policy; having an effective one in place may have prevented the initial vector of compromise in the first place.

Lessons learned
In the end, Ned learned to cover the basics, and then worry about being excellent (where have we heard that before?). Security practitioners like Ned must always ask simple questions of access and control. Sensitize employees to good security practices. Be vigilant—review network logs, patch vulnerabilities promptly, address the SQL injection issue—use two-factor authentication and have a strong password policy.

Top five victim-controllable investigative challenges
Over the years of conducting hundreds upon hundreds of investigations, we have seen our share of investigative challenges—those things that because they were not in place or were improperly configured, directly hampered or delayed an investigation, or required a deeper-dive analysis of other evidence sources. Our "top five victim-controllable investigative challenges" that we continue to come across are:

- Logs, logs, logs—specifically, the non-existence of, not enough of (rolling over too quickly) or difficulty in locating/retrieving in a timely manner.
- Network topologies—the lack of or the severely out of date.
- Baseline images and trusted task lists—the lack of, the inaccuracy or the out of date.
- "Dual-use" tools (for example, PsExec)—left on the system prior to its breach (and no, storing them in the Windows Recycler is not a security option), or no detection of their use.
- Self-inflicted anti-forensics—rebuilding systems and then calling us, containing and eradicating but not properly documenting, pulling the power cable and not the network cable, etc.

It goes without saying, having the appropriate amount (and type) of logging and being able to quickly identify and retrieve these logs goes a long way toward being able to quickly tackle all types of cybersecurity incidents. By the same token, having up-to-date and relevant network diagrams and data flow diagrams enables responders and investigators to have the ability to quickly scope, triage and envision the overall nature of the incident. Baseline images and trusted task lists can also provide responders and investigators with invaluable tools to separate the proverbial "wheat from the chaff" in terms of finding the badness.

Way forward

Well, this brings us to the end of our ride along. Let's take one quick look over your assembled Attack-Defend Card Deck.

Attack-Defend Card Deck	Takeaway
The human element	• Know the threat actors; recognize their methods. • Know your workforce; mitigate its vulnerabilities.
Conduit devices	• Know your environment; reduce its exposure.
Configuration exploitation	• Know your tools; configure properly.
Malicious software	• Know the threat actor tools; know their capabilities.

As with all strategies, these specific takeaways need to be tied together with two additional actions:

- Create a plan based on your people, your processes and your technology.
- Test and reassess your plan. Update as necessary.

Now it's time for you to take over the wheel. And in doing so, a few parting words of advice (and perhaps inspiration) from someone who intimately studied and experienced life in the "field."

"Know your enemy and know yourself and you can fight a hundred battles without disaster."—Sun Tzu

Questions? Comments? Brilliant ideas?
And be sure to give us your feedback to make this type of work product a more usable instrument in the future. Drop us a line at databreachdigest@verizon.com, find us on LinkedIn or on Twitter @VZdbir with the hashtag #VZreport.

Appendix A: Top 25 VERIS threat actions

The Top 25 VERIS Threat Actions for confirmed data breaches over the previous three years are as follows:[8]

1. Phishing – Phishing (or any type of *ishing)
2. Use of stolen creds – Use of stolen credentials
3. RAM scraper – RAM scraper or memory parser (capture data from volatile memory)
4. Brute force – Brute force attack
5. Export data – Export data to another site or system
6. Use of backdoor or C2 – Use of backdoor or C2
7. Unknown – Malware unknown
8. Backdoor – Backdoor (enable remote access)
9. Spyware/Keylogger – Spyware, keylogger or form-grabber (capture user input or activity)
10. Unknown – Hacking unknown
11. C2 – Command and control (C2)
12. Capture stored data – Capture data stored on system disk
13. Downloader – Downloader (pull updates or other malware)
14. Scan network – Scan or footprint network
15. Password dumper – Password dumper (extract credential hashes)
16. Privilege abuse – Abuse of system access privileges
17. Skimmer – Payment card skimmers
18. Adminware – System or network utilities (e.g., PsTools, Netcat)
19. Rootkit – Rootkit (maintain local privileges and stealth)
20. SQL injection – SQL injection attack
21. Exploit vuln – Exploit vulnerability in code (vs misconfig or weakness)
22. Disable controls – Disable or interfere with security controls
23. Brute force – Brute force attack
24. Unapproved hardware – Use of unapproved hardware or devices
25. Packet sniffer – Packet sniffer (capture data from network)

8 http://veriscommunity.net/enums.html#section-actions

Appendix B: CIS critical security controls

The 20 Center for Internet Security (CIS) critical security controls (CSCs)—Version 6 are as follows:[9]

CSC 1	Inventory of Authorized and Unauthorized Devices
CSC 2	Inventory of Authorized and Unauthorized Software
CSC 3	Secure Configurations for Hardware and Software
CSC 4	Continuous Vulnerability Assessment and Remediation
CSC 5	Controlled Use of Administrative Privileges
CSC 6	Maintenance, Monitoring, and Analysis of Audit Logs
CSC 7	Email and Web Browser Protections
CSC 8	Malware Defenses
CSC 9	Limitation and Control of Network Ports
CSC 10	Data Recovery Capability
CSC 11	Secure Configurations for Network Devices
CSC 12	Boundary Defense
CSC 13	Data Protection
CSC 14	Controlled Access Based on the Need to Know
CSC 15	Wireless Access Control
CSC 16	Account Monitoring and Control
CSC 17	Security Skills Assessment and Appropriate Training to Fill Gaps
CSC 18	Application Software Security
CSC 19	Incident Response and Management
CSC 20	Penetration Tests and Red Team Exercises

9 https://www.cisecurity.org/critical-controls.cfm

verizonenterprise.com

© 2016 Verizon. All rights reserved. The Verizon name and logo and all other names, logos and slogans identifying Verizon's products and services are trademarks and service marks or registered trademarks and service marks of Verizon Trademark Services LLC or its affiliates in the United States and/or other countries. All other trademarks and service marks are the property of their respective owners. WP16641 2/16