Attack-Defend Card

HE-2: Hacktivist Attack – the Epluribus Enum

Breach scenario

Breach scenario
Specific

Sophistication level
1 — **2** — **3** — 4 — 5

Attributes
Confidentiality, Integrity

Incident pattern

Pattern
Web application attacks, DoS attacks

Time to discovery
H — **D** — W — M — Y

Time to containment
H — **D** — **W** — M — Y

Threat actor

Composition
Activist

Motives
Ideology, Grudge

Tactics and techniques
DoS, Unknown hacking, Backdoor, Use of backdoor or C2, C2

Targeted victims

Industries
Financial, Public, Information

Key stakeholders
Legal Counsel, Corporate Communications, Incident Commander

Countermeasures
CSC-4, CSC-5, CSC-7, CSC-16, CSC-18

Description
Hacktivist attacks leverage hacking techniques as a form of activism; these differ from the multitude of other attack types by the unique motivation of the threat actor. Commonly a hacktivist is motivated by a desire to harm or embarrass their targeted victim in an effort to further a political or social agenda.

20

An Executive Doxxing Match

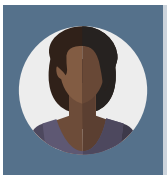

Stakeholder

Lead Investigator

The situation

The Verizon RISK Team was contacted by a multinational organization that had attracted negative attention following the handling of an unpopular company restructuring. This customer, Cheese Movers International (CMI), had a significant number of disgruntled employees and ex-employees. They had also drawn the attention of more than one group of hacktivists who had posted messages on their social media accounts referencing the changes. Various derogatory hashtags on social media were popping up and threats against executives were being posted to social networking sites.

The customer was a soft target for hacktivism; their attack surface was large due to their sheer size and their diverse, global business units. This was exacerbated further by the risk of an insider threat or recently terminated ex-employee using their advanced knowledge of the organization to perpetrate an attack or to leak information assisting other threat actors.

On the face of it, there was no evidence that any attack had been initiated; however, CMI sought our assistance to help them proactively gather threat intelligence, perform penetration testing, and be prepared should any of the online threats materialize.

Response and investigation

We initially provided CMI with assistance and guidance in collating and reviewing open-source intelligence; this included searching social networks and online forums as well as specialized investigative activities within the darknet, the less accessible part of the internet, which is anonymized by protective software and configurations. We set up a secure anonymous account of our own, which enabled us to search through marketplaces and other locations on the darknet to see what the hacktivists were discussing in relation to CMI. These activities identified a huge number of threats and negative statements. And although the majority was not considered genuine, the home address and personal details of executives were being actively sought by suspicious parties.

> The customer was a soft target for hacktivism; their attack surface was large due to their sheer size and their diverse, global business units.

Later on, evidence was found that personal details for two executives had been obtained and were being shared online. CMI was able to implement the Incident Response (IR) Plan they had developed to deal with this type of situation as it arose. The breach of personal information associated with senior executives was identified early enough that it could be reported to Law Enforcement (LE) before malicious parties acted upon it; as a result, the ensuing threatening phone calls and spurious deliveries were monitored from the outset and were immediately followed up. Local LE also provided a liaison officer and guidance on physical protection considering the threats that had been received.

Unfortunately, this was just the first of multiple threats and attacks experienced over the course of the next three weeks. Distributed Denial of Service (DDoS) attacks were attempted against many of the company's websites. The majority of these were thwarted by the DDoS protection capability that CMI had put in place as a result of the intelligence provided by our Verizon Cyber Intelligence Center (VCIC) and our Darknet Research Team.

We collaborated with our Pen Testing Team to perform urgent assessments of key assets. Due to the very short timeframe, these assessments were performed on a best effort basis, but they successfully identified vulnerabilities in web-facing servers which could have proven catastrophic had they been noticed by hacktivists. In two cases, a Structured Query Language (SQL) injection vulnerability and an unpatched application with known vulnerabilities were identified. It was later found that both servers had been targeted with reconnaissance activities, which may have identified the same vulnerabilities had they not been urgently patched by CMI.

After approximately two weeks of successfully defending against attacks on all fronts, an attack was finally successful. One of CMI's websites appeared to have been defaced: The site was not accessible and had been replaced with a message claiming responsibility and blaming CMI for inviting this retribution. The posted message claimed that CMI servers had been hacked and customer data would be leaked unless certain actions were performed. We quickly determined the defacement did not appear to be the result of a compromised CMI system, but rather visitors to the relevant Uniform Resource Locator (URL) were being redirected to another server hosting the message.

As a matter of due diligence, we deployed our investigators to the datacenter containing the affected web server. We quickly confirmed that no evidence of a breach existed. Furthermore, our RISK Network Forensics (NetFor) Team, who had previously deployed full packet capture devices within four data centers, had not identified any suspicious or malicious traffic.

It was later determined that the domain registrar for the effected domain had been targeted in a social engineering attack, during which the threat actor successfully impersonated CMI staff. They were able to gain access to the account on the domain registrar's service and modify the relevant Domain Name System (DNS) records, which caused visitors to the CMI URL to be redirected to another website.

Fortunately, the site in question was not CMI's principal website and was only used by a small subset of their customers. The DNS issue was quickly resolved and eventually this domain was migrated to their principal domain registrar, whose security practices were superior.

As with many similar incidents the media attention soon dried up as did the interest of the hacktivists. The DDoS attacks became less and less frequent and the internet was soon engulfed in the next drama. CMI maintained extra vigilance for a number of months, but before long, it was back to business as usual.

Lessons learned

The information gathered from intelligence sources was vital in our response efforts, as it provided us with the knowledge of who the targeted victims were, and the tactics the threat actors would deploy. Mitigation and response activities are as follows:

Mitigation

- **Don't rock the boat.** Stay off the radar of any potential hacker.
- **Keep an ear to the ground.** Base defenses, detection mechanisms and response capabilities on sound threat intelligence.
- **Secure your environment.** Implement a timely and effective patch management program; conduct regular penetration-testing activities.
- **Protect social media accounts.** Use two-factor authentication, strong and varied passwords, as well as proper security awareness training for staffs who manage the social media presence.
- **Protect third-party services.** Protect account credentials; use a reputable domain name registrar that offers two-factor authentication or approved IP address whitelisting[7].

Response

- **Prepare and initiate your IR Plan.** Establish an IR Plan early, and then regularly review, test and update it.
- **Scope and triage the incident quickly.** Effectively scope and task prioritize; be prepared to manage simultaneous, yet distinct, incidents.
- **Proactively communicate with affected entities.** Confirm facts quickly; develop a remediation strategy and communicate this to customers.
- **Engage LE at the right time.** Consider legal and regulatory responsibilities in conjunction with advice from Legal Counsel; contact LE when the time is right.

7. For further details on the recommendations above, see "Data Breach Digest – October 2016 Update, Hacktivist Attack: Shedding Light on the Matter" at www.verizonenterprise.com/verizon-insights-lab/data-breach-digest/2016/.

Attack-Defend Card

HE-3: Partner Misuse – the Indignant Mole

Breach scenario

Breach scenario
Specific

Sophistication level
1 — 2 — 3 — 4 — 5

Attributes
Confidentiality, Integrity

Incident pattern

Pattern
Insider and privilege misuse

Time to discovery
H — D — W — M — Y

Time to containment
H — D — W — M — Y

Threat actor

Composition
Other (Partner)

Motives
Financial, Espionage, Grudge

Tactics and techniques
Data mishandling, Net misuse, Privilege abuse

Targeted victim

Industries
Accommodation, Financial, Retail, Healthcare

Key stakeholders
Legal Counsel, Incident Commander, Corporate Communications

Countermeasures
CSC-6, CSC-12, CSC-13, CSC-16, CSC-19

Description

Partner misuse involves semi-trusted entities who have some level of enterprise environment access and, either through purposeful maliciousness or inadvertent ineptitude, lead to a breach of that environment.

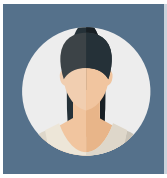

Stakeholder

Legal Counsel

The Broken Circle of Trust

The Situation

As a partner in a global law firm, I have 25 years of experience assisting clients who manage various litigation and employment issues resulting from data breaches. Without necessarily realizing it, I've been working in the "data breach" industry for many years, be it advising my clients on internal investigations and associated litigation involving cybersecurity and privacy issues, or helping them meet their regulatory requirements following inadvertent data disclosures.

> [The company was] seeking my advice regarding their obligations concerning data protection and other relevant cybersecurity legislation.

Recently one of our clients, a regional water supplier, contacted my firm to discuss an incident that affected several of their small and medium-sized enterprise clients. Their clients had recently notified them that their online account details had changed. The company had wisely identified the potential that customer data had been compromised and they were seeking my advice regarding their obligations concerning data protection and other relevant cybersecurity legislation.

Unfortunately, as our conversation progressed, the issue extended beyond a simple data breach. When customers had their passwords reset and regained access to their accounts, many noticed that the registered bank account details had also been changed. This meant that refunds due to the customers had been transferred fraudulently to new bank accounts. It was later determined that the refunds totaled over £500,000 and were directed to two bank accounts in England.

I subsequently worked with Law Enforcement (LE) and the National Action Fraud Hotline to track down the bank account holder. As I did so, it became clear that the banks had also been socially engineered. Believing the refunds to be foreign deposits, they allowed the account holder to transfer 90% of the money to accounts in Dubai and the Bahamas as soon as the payments arrived in their UK account. Ultimately, the funds had been withdrawn from the accounts and used to purchase Bitcoin, which was transferred to addresses associated with a Bitcoin mixing (laundering) service. The trail went cold and the LE inquiries failed to identify a subject.

After several discussions with my client, it was obvious they had had a data breach. What wasn't obvious was how the breach occurred. Despite a robust security posture, none of their security appliances or log sources showed any signs of compromise. A review of affected accounts and systems showed no signs of malware or tampering. With my client's approval, I reached out to the Verizon RISK Team for assistance in the investigation – hoping desperately that they could turn up some new evidence.

Threat targeting: Attention small business owners!

Small business owners often consider themselves as an unlikely target, believing themselves to be a smaller fish in the sea for an attacker. However, this feeling evidently provides a false sense of security when we consider the high amount of secondary attacks conducted from compromised systems.

Attackers often compromise smaller less secure businesses and use their environments as their base of operations. The attackers rely on relatively insecure systems with poor monitoring and logging as an additional layer of security when perpetrating attacks. Your systems might be the origin of major breaches and, in addition, your intellectual property might be an attractive bonus.

Response and investigation

Once the RISK Team arrived at my client's premises, a "war room" was established and the discussions turned to network diagrams, web servers, log files and payment and refund flows. While some of the technical details went over my head, the RISK Team hit the proverbial ground running. They were quickly able to establish all the systems and processes involved in managing the customer account creation and storage.

> A "war room" was established and the discussions turned to network diagrams, web servers, log files and payment and refund flows.

The RISK Team did a due diligence review of various logs and the web server itself. Using their listing of known Indicators of Compromise (IoCs) they confirmed that no malicious software was present. With very little to go on from a technical standpoint, the RISK Team lead investigator suggested we speak with some of the people involved. I expressed my concerns – it would be a large project to interview so many people, and many employees were remotely located in India. The RISK Team investigator assured me this was nothing new for his team and he already had resources lined up in India ready to travel onsite.

Agreeing to the plan, my customer allowed the RISK Team to conduct interviews with various stakeholders including those identified at a third-party call center in Mumbai, India. This call center was responsible for administering the online accounts and processing telephone payments. Two RISK Team investigators arrived in Mumbai to interview the third-party call center personnel. During the interview, and subsequent review of the Customer Relationship Management (CRM) log files, it became evident that one user had accessed all the accounts that had been fraudulently refunded.

An investigation of the user's computer confirmed the access to my client's Content Management System (CMS) records in question; however, there was nothing to suggest the data had been copied or that the refunds had been requested using this computer. The user denied any knowledge of the fraudulent activity and suggested the computer must have been hacked, although the RISK Team's analysis identified no such evidence. The user was so adamant that he was not involved that to "prove" it he signed an affidavit that permitted the RISK Team to examine his home computer.

Security imperative: Multi-factor authentication

Multi-Factor Authentication (MFA) is an access control system that allows users to authenticate to resources using two or more independent forms of identification. These fall into the categories of something you know, such as a user-created password, something you have, such as a one-time passcode (OTP), and someone you are, such as your fingerprint or retina scan. A unique OTP is typically generated every 60-90 seconds on a physical dongle or within an application. This requires physical possession to be read (thereby aligning with the "something you have" factor). Many users may have an application installed on their smartphone through which they can obtain the OTP at any time.

When a user authenticates to an MFA system, the system first checks that it has received the correct user-created password (something you know). Next it checks the OTP (something you have), which is known only to the user and the MFA system. Only when these two pieces of information are correct does the MFA system allow the user to authenticate successfully.

An alternative method of MFA involves a known password and a biometric scan for authentication. Using this method, a user may authenticate by providing a user-created password (something you know) in addition to a biometric scan (something you are). These biometric scans are typically done on a user's fingerprint or retina, as these are unique to each individual. Additionally, many hardware and software developers have started to introduce facial recognition technology as a means of biometric authentication.

An initial review of the user's home computer system revealed very little data. In fact, so little was found on the system that it appeared to have been systematically cleaned using data wiping software. Unfortunately, for the user, the wiping software did not fully clean the volume. Shadow copies of data were recovered revealing numerous email messages between the call center employee and another individual, later identified to be his cousin in the UK. These emails contained pictures of account details that correlated to the accounts affected by the fraudulent activity. The RISK Team pointed out that the metadata within these photos indicated that they had been taken with a camera phone and the photos appeared to be of a computer system monitor.

> Shadow copies of data were recovered revealing numerous email messages between the call center employee and another individual, later identified to be his cousin in the UK.

With new evidence in hand, the RISK Team returned to the Mumbai office for a follow-up interview with the suspected worker. When presented with the data retrieved from his home computer, the worker finally confessed to the crime and offered assistance in identifying accounts with over £1,000 in refunds stolen.

Working with LE and the RISK Team, a plan was hatched to verify the identity of the employee's cousin. The employee would take a photograph of the account details and would send the picture to his cousin in England, who would then create an online account or request a password reset for their current account as he had in the past. Once we validated the change was in place, we took the phone number and log file evidence to the authorities to secure a conviction.

Lessons learned

It's always good to sit back, relax, and reflect after an incident. The main points of consideration coming out of this incident would be to review in-place agreements with partners who have access to your critical data and that they conduct stringent background checks on their employees. Typical mitigation and response actions to take for partner misuse situations are:

Mitigation

- Monitor corporate and guest network activity.
- Take steps to reduce external device threats.
- Keep tabs on sensitive data.
- Be cognizant of changes in employee attitude/behavior.
- Establish a Data Classification Policy (and limit printing copies).

Response

- Prepare and initiate your IR Plan in a timely manner.
- Quickly scope and triage the incident.
- Proactively communicate with affected entities.
- Seek advice from Legal Counsel; contact LE when the time is right.

Attack-Defend Card

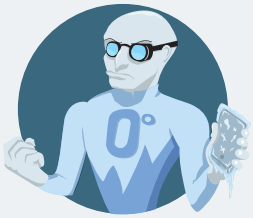

HE-4: Disgruntled Employee – the Absolute Zero

Breach scenario

Breach scenario
Specific

Sophistication level
1 — 2 — 3 — 4 — 5

Attributes
Confidentiality, Integrity

Incident pattern

Pattern
Insider and privilege misuse

Time to discovery
H — D — W — M — Y

Time to containment
H — D — W — M — Y

Threat actor

Composition
Other (Employee)

Motives
Grudge, Espionage

Tactics and techniques
Export data, Privilege abuse, Capture stored data, Disable controls

Targeted victim

Industries
Public, Financial, Healthcare

Key stakeholders
Human Resources, Legal Counsel, Incident Commander

Countermeasures
CSC-1, CSC-6, CSC-10, CSC-13, CSC-16

Description

Disgruntled employees, especially those disillusioned with their company, can represent one of the most difficult threat actors against which to defend. Layoffs, pay cuts, or organizational shifts may leave some employees in a position where they can rationalize nefarious activities.

A "Pre-Competitive" Advantage

Stakeholder

Human Resources

The Situation

By definition, employees have access to privileged systems and information; this means large amounts of legitimate activity will need to be sorted through during breach response efforts. Any employee can be angry enough to do something malicious, and therefore special care needs to be taken around events that can increase employee emotions.

Firing people was rarely an interesting job, but as I sat filling out the final forms for terminating Mr. Simpson, I breathed a sigh of relief, glad to be done with the ordeal. On the surface, it seemed like a straightforward case. Mr. Simpson's team was being merged with another team and he was unhappy with the new hierarchy. After being informed by a friend in Human Resources (HR) about the upcoming changes, Mr. Simpson began using his administrative access to take over other accounts. He ultimately attempted to disrupt operations – a vindictive response to being underappreciated – and downloaded confidential files (a bargaining chip for this next job). It seemed so cut and dried – he did it and admitted to it – but still the lawyers required us to collect the evidence to prove it.

Response and investigation

I don't imagine most investigations begin with the answer, but with a very candid confession from the primary suspect, ours did. We knew how, when, and what happened from Mr. Simpson's description and by the time we engaged the Verizon RISK Team, all we needed them to do was document and verify the claims from a technical point of view. Once we knew we had the whole truth, I could then expect to fill out a stack of forms to safely terminate Mr. Simpson's employment.

The events that led to Mr. Simpson's confession were well-documented. On an otherwise normal Friday afternoon, a programmer reported that an application was experiencing unexpected failures and an internal investigation began. This investigation turned up multiple suspicious log entries showing Mr. Simpson logging into the application server only minutes before the problems started. The logs showed failed super user account access from Mr. Simpson, followed by password resets of service accounts. These findings could potentially have been legitimate as Mr. Simpson was an IT administrator, but the circumstances surrounding them – no ticket or prior notification – led to the interviews in which he eventually revealed his actions, in hopes of leniency.

Incident pattern: Insider and privilege misuse

The "Privilege Misuse" pattern is one of the few that includes collusion between internal and external actors. According to VERIS, the top industries affected by this are the public sector, healthcare, and finance organizations. This category covers the insider threat, but can also include external actors collaborating with internal actors to gain unapproved or malicious use of organizational resources.

Financial gain and espionage remain the primary motivation for committing this type of attack. The most common form of misuse is merely using access to gain information for alternative and unsanctioned uses. The weakest link for any organization is not its systems, but rather the human factor. It is important to note that these incidents are not always the result of a malicious employee and often stem from carelessness and lack of awareness regarding sound IT protocol.

Insider threats are usually the most difficult to detect and can take months, or longer, to discover. Identifying insider privilege abuse can be difficult because it is often committed by employees perceived to be trustworthy, and because they are using the privileges granted to them by the organization. Organizations should proactively take steps to minimize the privileges users are provided with. They should also keep detailed audit logs of users with administrative privileges.

In addition to the known application server activities, Mr. Simpson admitted to accessing multiple email boxes using the service accounts to collect data for interview use and to insert scheduled jobs designed to disrupt his new team's workflows. This was a lot of data to sort through, and I honestly didn't know where to begin looking to verify these claims. Thankfully, our IT Security Department had called in the RISK Team to assist in the digital forensic examination to determine if Mr. Simpson had left any other surprises for us to find.

The RISK Team requested a huge number of log files and mailbox summaries, and immediately started digging in. It was only the next day when preliminary findings began coming back to us. The investigators verified that Mr. Simpson had used his access to compromise other accounts. Much to my surprise, included in their initial findings was a listing of every file he had downloaded from another user's inbox, which looked like it included everything from operations documents to product technical details. This was more than a bargaining chip. This was corporate theft. Beyond the stolen files was a second listing of scheduled jobs inserted by Mr. Simpson. The jobs were exclusively mass delete commands scheduled to occur at critical times over the next year: During tax season, prior to holiday bonuses, and a few seemingly random dates.

> The jobs were exclusively mass delete commands scheduled to occur at critical times over the next year: During tax season, prior to holiday bonuses, and a few seemingly random dates.

While our internal teams worked to remove the jobs and validate the contents of each stolen file, the RISK Team investigators moved on to their second phase - discovering any other activity to which Mr. Simpson may not have confessed. After requesting "network logs" from our IT Security Team, the investigators turned to searching for known threat actors and suspicious activity. They also focused analysis on the time range defined by the service account compromises. A few days and a dozen email requests later, a second set of findings arrived from the RISK Team.

The RISK Team review of the network traffic had identified suspicious connections to a server in Romania. This particular server was owned by a short-term lease hosting location using Bitcoin as payment. The report explained that this was currency used frequently by hackers wishing to remain anonymous, and while completely unrelated to Mr. Simpson's activity, many other attacks had involved this system. Closing out the findings was a set of instructions for our IT Security Team on how to find and identify the internal system in question.

> [Bitcoin] was currency used frequently by hackers wishing to remain anonymous, and while completely unrelated to Mr. Simpson's activity, many other attacks had involved this system.

It took our IT Security Team only a few hours to find the suspicious system and remove it from the network for further review. The onsite RISK Team investigators collected a forensic system image and shipped it to the RISK Labs for examination. This proved fruitless; comparisons with known malicious files and analysis of changes around the time of the network activity revealed nothing. Both the IT Security Team and RISK Team were baffled, as the traffic was definitely coming from this system and had stopped immediately after the device was taken offline; however, nothing seemed to be out of place. We were getting antsy.

Returning to the physical device, the RISK Team investigators began to collect additional forensic information and had a lucky break. While plugging in a USB keyboard to issue commands, the investigator noticed an extension on the plug itself. When pried, it popped off, revealing an off-the-shelf, clandestine keylogger. The RISK Team explained that the keylogger was designed to capture any input a user provided via the keyboard and was sending the capture to the rented Romanian server. I was stunned; this was the kind of thing I thought I'd see in a movie, not my job, but the proof was there in our hands.

Mr. Simpson's actions were vindictive and done in response to the recent restructuring of the company's IT Department. One of Mr. Simpson's main motivations was to make the new IT Department appear incompetent. He had admitted that he was planning to use the information he stole as leverage in finding a job with a competitor and possibly profit from his exploits. Finally, he had lied about the extent of his actions and clearly had gone beyond simply being upset. With the evidence and paperwork in hand, Mr. Simpson was summarily fired and the forensic reports were provided to law enforcement.

> While plugging in a USB keyboard to issue commands, the investigator noticed an extension on the plug itself. When pried, it popped off, revealing an off-the-shelf, clandestine keylogger.

Lessons learned

Our company narrowly dodged a bullet in that some of our most sensitive information and intellectual property was nearly stolen; we learned several lessons as a result of this incident. One lesson was that the friend in HR should not have notified Mr. Simpson. Details regarding restructuring and moving of specific jobs should be closely held and carefully coordinated with department managers. Another was the company should have had an action plan in place, such as increased monitoring of employees affected by the transition, to reduce the risk of vindictive behavior by those affected. Finally, as part of the transition, the company should also have conducted a thorough asset inventory. Doing so might have identified any installed keyloggers.

Actionable intelligence: No, it's not an oxymoron

Some say the two words "actionable intelligence" together form an oxymoron. However, in cybersecurity parlance it has a very critical objective by the virtue of being (a) actionable; i.e., a clear set of actions or countermeasures that can help prevent or detect a cyberattack and (b) intelligence; i.e., sometimes abstract but relevant knowledge that can help pinpoint a cybersecurity event, trend, pattern or incident from the perspective of known-unknowns as well as unknown-unknowns.

Actionable intelligence is derived from telltale signs or early threat warnings based on external information or feeds, and internal threat hunting efforts. This should ideally be specific enough to identify compromised assets, usable as an indicator of compromise and, most importantly, consumable by prevention and detection security platforms. There has to be a clear difference between "raw information" or "information overload" and the real "actionable intelligence." Organizations should try to elevate their focus from having access to indicators of compromise to indicators of anomalies.

Attack-Defend Card

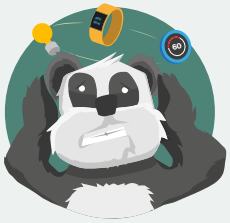

CD-3: IoT Calamity – the Panda Monium

Breach scenario

Breach scenario
Opportunistic (IoT Devices), Indirect (DoS Attack Victim)

Sophistication level
1 — **2** — **3** — 4 — 5

Attributes
Availability

Incident pattern

Pattern
DoS attacks, Insider and privilege misuse, Crimeware

Time to discovery
H — D — W — M — Y

Time to containment
H — D — W — M — Y

Threat actor

Composition
Activist, State-affiliated

Motives
Grudge, Ideology, Financial

Tactics and techniques
Brute force, Privilege abuse, Scan network, Exploit vulnerability

Targeted victims

Industries
Entertainment, Professional, Educational, Administrative, Information, Manufacturing

Key stakeholders
Incident Commander, Legal Counsel, Corporate Communications

Countermeasures
CSC-1, CSC-3, CSC-9, CSC-11, CSC-12

Description
Security is often an afterthought when it comes to Internet of Things (IoT) solutions – and that means devices are often vulnerable to a wide array of threats. IoT calamity attacks take advantage of these cybersecurity shortfalls in IoT devices.

A Botnet Barrage

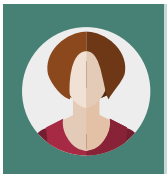

Stakeholder

Incident Commander

The situation

Senior members of my university's IT Security Team rotated weekly as on-call "Incident Commanders" in the event that a response was needed. This week was my turn and as I sat at home, my phone lit up with a call from the help desk. They had been receiving an increasing number of complaints from students across campus about slow or inaccessible network connectivity. As always seemed to happen, the help desk had written off earlier complaints and it was well after 9 PM when I was finally pulled in.

I joined the conference bridge and began triaging the information. Even with limited access, the help desk had found a number of concerns. The name servers, responsible for Domain Name System (DNS) lookups, were producing high-volume alerts and showed an abnormal number of subdomains related to seafood. As the servers struggled to keep up, legitimate lookups were being dropped – preventing access to the majority of the internet. While this explained the "slow network" issues, it raised much more concerning questions. From where were these unusual DNS lookups coming? And why were there so many of them? Were students suddenly interested in seafood dinners? Unlikely. Suspecting the worst, I put on a pot of coffee and got to work.

Response and investigation

Now that I had a handle on the incident in general, I reached out to the Verizon RISK Team, who we had on retainer, and began the process of escalating the issue. At their request, I gathered up network and firewall logs and passed them along for review. My IT Security Manager assured me that review would begin immediately and listed off a few of the triage steps he would be taking. All logs would be processed for known indicators of malicious activity and firewall logs in particular would be used to identify the sources of these requests.

Within hours, I had more feedback than I could handle and began the review process. The firewall analysis identified over 5,000 discrete systems making hundreds of DNS lookups every 15 minutes. Of these, nearly all systems were found to be living on the segment of the network dedicated to our IoT infrastructure. With a massive campus to monitor, everything from light bulbs to vending machines had been connected to the network for ease of management and improved efficiencies. While these IoT systems were supposed to be isolated from the rest of the network, it was clear that they were all configured to use DNS servers in a different subnet.

The impact of IoT

IoT possesses a huge potential to forever change the way we interact with our world through technology. The proliferation of IoT devices essentially leads to increased automation, big data analytics, and artificial-intelligence-based decision making in our daily lives. An IoT solution requires a detailed and comprehensive security and privacy framework – an area that unfortunately still requires a lot of work on design – as well as a substantial impetus on collaboration by the IoT market players on the underlying security.

Despite the fact that we are in a hyper-connected world, the security of the IoT is still at times somewhat of an afterthought. The main issue is that most firms do not realize that components behind the IoT's agile innovation can easily go wrong, and can have a far greater impact than what can be seen in the traditional IT landscape. IoT devices are usually constantly connected to the internet and may not be looked at from a security perspective, thus leaving them vulnerable to a variety of attacks. This makes IoT devices an ideal target for being conscripted into a botnet army.

The RISK Team provided me with a report detailing known indicators found in the firewall and DNS logs that I had sent over earlier. Of the thousands of domains requested, only 15 distinct IP addresses were returned. Four of these IP addresses and close to 100 of the domains appeared in recent indicator lists for an emergent IoT botnet. This botnet spread from device to device by brute-forcing default and weak passwords. Once the password was known, the malware had full control of the device and would check in with command infrastructure for updates and change the device's password – locking us out of the 5,000 systems.

This was a mess. Short of replacing every soda machine and lamp post, I was at a loss as to how to remediate the situation. We had known repeatable processes and procedures for replacing infrastructure and application servers, but nothing for an IoT outbreak.

> Short of replacing every soda machine and lamp post, I was at a loss as to how to remediate the situation.

Luckily, for me, a less drastic option existed than replacing all the IoT devices on campus. Analysis of previous malware samples had shown that the control password, used to issue commands to infected systems, was also used as the newly updated device password. These commands were typically received via Hypertext Transfer Protocol (HTTP) and in many cases did not rely on Secure Sockets Layer (SSL) to encrypt the transmissions. If this was the case for our compromise, a full packet capture device could be used to inspect the network traffic and identify the new device password. The plan was to intercept the clear text password for a compromised IoT device over the wire and then use that information to perform a password change before the next malware update. If conducted properly and quickly, we could regain control of our IoT devices.

While we waited for the full packet capture solution to be set up, I instructed the Network Operations Team to prepare to shut down all network access for our IoT segments once we had intercepted the malware password. Short lived as it was, the impact from severing all of our IoT devices from the internet during that brief period was noticeable across the campus – and we were determined never to have a repeat incident.

> The plan was to intercept the clear-text password for a compromised IoT device over the wire and then use that information to perform a password change before the next malware update.

Lessons learned

With the packet capture device operational, it was only a matter of hours before we had a complete listing of new passwords assigned to devices. With these passwords, one of our developers was able to write a script, which allowed us to log in, update the password, and remove the infection across all devices at once. The whole process took a matter of minutes and I made a mental note to save that script for later – although I prayed that we would never need it again. Now that the incident had been contained, we looked toward ways to prevent it from happening again.

Mitigation

- Don't keep all your eggs in one basket; create separate network zones for IoT systems; air gap them from other critical networks where possible.
- Don't allow direct ingress or egress connectivity to the internet; don't forget the importance of an in-line proxy or content filtering system.
- Change default credentials on devices; use strong and unique passwords for device accounts and Wi-Fi networks.
- Regularly monitor events and logs; hunt for threats at endpoints, as well as at the network level; scan for open remote access protocols on your network and disable commonly unused and unsecured features and services (such as Universal Plug and Play (UPnP) and Real Time Streaming Protocol (RTSP)) that aren't required.
- Include IoT devices in IT asset inventory; regularly check manufacturer websites for firmware updates.

Response

- Develop and follow your pre-designed IR playbooks to tackle IoT device-related incidents.
- Scope and contain incident immediately; segregate affected subnet and restrict network ingress and egress communication to/from affected subnet.
- Change admin or console passwords of the IoT systems and controllers.
- Leverage network forensics, to include network logs, NetFlow data, and packet captures.
- Consider informing Law Enforcement (LE) and regional Computer Emergency Readiness Team (CERT) organizations as egress communication may have impacted other entities and the related threat intelligence could help other potential victims.

The evolution of the IoT

Like any typical Gen-X technology, the IoT continues to evolve and has gone through a growth spurt over the past few years. This rapid proliferation has led to as many new issues as the underlying devices were intended to solve.

The underlying problem is that many IoT manufacturers are primarily designing their devices for functionality; and proper security testing often takes a back seat. It's even more necessary with IoT devices that the buyer scrutinizes the security of any devices they use. IoT botnets spread quickly because they don't face some of the problems conventional botnets do, due to the fact that IoT devices are often rarely patched or updated.

In addition, the vendors that create IoT devices, along with the users that own and operate them, aren't always directly impacted by a compromise or even immediately aware that their devices played a role in a cybersecurity incident. In a number of these circumstances, the IoT environment leveraged in an attack is not actually the intended victim, but rather an involuntary accomplice that is being used to attack an unrelated third-party target.

IoT threats go well beyond a typical security breach where concerns revolve around the theft of confidential data. In this new age of IoT breaches, we are seeing a growing and wide-ranging impact in our physical world as well as on human life/safety (e.g., transportation or medical device incidents) and even a changing financial and legal liability landscape.

Today, the IoT is not confined within an organization's typical control boundary, as the connected infrastructure has moved far beyond those control lines. These devices exist virtually everywhere, are available anytime, and are on a variety of platforms. This must prompt organizations to think about IoT threat modeling in a manner that incorporates security and privacy by design.

Attack-Defend Card

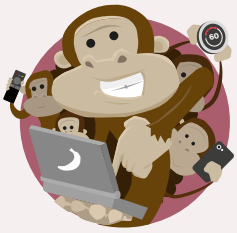

CE-2: DDoS attack – the 12000 Monkeyz

Breach scenario

Breach scenario
Specific

Sophistication level
1 — **2** — **3** — 4 — 5

Attributes
Availability

Incident pattern

Pattern
DoS attacks

Time to discovery
H — D — W — M — Y

Time to containment
H — **D** — W — M — Y

Threat actor

Composition
Activist, State-affiliated

Motives
Grudge, Ideology, Financial

Tactics and techniques
Brute force, Privilege abuse, Scan network, Exploit vulnerability

Targeted victims

Industries
Entertainment, Professional, Educational, Administrative, Information, Manufacturing, Retail

Key stakeholders
Incident Commander, Corporate Communications, Legal Counsel

Countermeasures
CSC-3, CSC-9, CSC-11, CSC-12, CSC-19

Description

A Denial of Service (DoS) attack involves a single computer using its network connection to flood a targeted system or resource with traffic. Distributed Denial of Service (DDoS) attacks leverage large numbers of systems to disrupt network operations across large networks.

No Patch, No Service

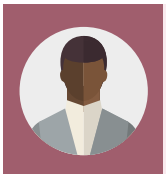

Stakeholder

SOC Analyst

The situation

DDoS attacks seem to be climbing at a steady rate year over year. The motivations for such attacks range from disrupting hostile competition, extortion, and political objectives. Although the incentive to launch a DDoS is rarely exfiltration of data, disruptions of a service or product can be just as devastating for any business. With the rise in popularity of DDoS attacks for threat actors, toolkits to launch these attacks have become easier to use and more effective by increasing overall bandwidth capabilities. Preparations for a DoS or DDoS attack include having the right team to handle the situation and is a critical component of the mitigation and recovery phases when dealing with these types of attacks.

> With the rise in popularity of DDoS attacks for threat actors, toolkits to launch these attacks have become easier to use and more efective by increasing overall bandwidth capabilities.

As a Security Operations Center (SOC) analyst, the ability to leverage tools and resources – in-house, external, or social media – definitely helps defend against some of the most aggressive attacks during pivotal times for the business.

During one of the largest volumetric attacks against a company in the software-as-a-service sector, I stood in the front lines of an uphill battle that exhausted all response team resources. Ultimately, this event shaped the way product launches and security were handled in the future for the company.

We determined the objective of the threat actor was solely to disrupt a holiday week and, in doing so, deny clients access to tools essential to handling their holiday workload. This well-timed attack coincided with a new product release date and a week in which a substantial influx of users was expected. Thus, the attack against our bandwidth would be compounded with tens of thousands of legitimate users trying to connect simultaneously.

With an attack of such great magnitude, the identifiers came in various forms – NetFlow graphs showed a 300 percent increase in the sample; Top Talkers lit up the target prefix to which most of the traffic was destined; and point-to-point protocol (PPP) Generic Routing Encapsulation (GRE) tunnels started to bounce up and down due to oversaturation. As a result, some applications were inaccessible to users wishing to access their accounts.

Our capability to view network traffic live with packet analysis tools played a major role in the active mitigation process. Review of the collected packets revealed four distinct types of DDoS: A Simple Service Discovery Protocol (SSDP) flood; a SYN flood; a Transmission Control Protocol (TCP) flood using invalid flag combinations; and a User Datagram Protocol (UDP) flood to non-web ports.

With so many types of DDoS, the priority for me as a SOC analyst was to mitigate what I could and attempt to recover the systems to a usable state. While the rest of the SOC and I worked to deal with the issue, the Verizon RISK Team was tracking the source of the threat actors, and investigating the extent of the threat actor's actions.

Response and investigation

One of the major challenges my organization dealt with when responding to the attack was routing the flood to our DDoS mitigation provider. When the DDoS attack occurred, we found our IT team underprepared and unable to quickly adjust our publicly advertised border routes. Initially, the routes were added to pass traffic through a scrubbing service prior to being sent to our servers; however, without clear documentation the engineer making the changes left the existing routes in place. This small oversight allowed roughly half of the incoming traffic to bypass the DDoS mitigation provider. After a tense hour diagnosing the problem, we discovered and corrected the error allowing us to move on to other forms of mitigation.

> Most of these systems were compromised routers running old firmware with UPnP enabled; odds were that many of these were “NYP’d” (not yet patched).

With the proper routing in place, we were able to begin handling the discrete attacks. Source-to-destination Access Control Entries (ACEs) were used to mitigate most of the SSDP and Invalid TCP flag combinations. This single action reduced a large portion of the attack traffic; however, considering the overall size of this attack there was still work to be done.

The non-spoofed IP addresses were reviewed and it was revealed that each had an open SSDP port (1900), which was publically accessible from the internet. Most of these systems were compromised routers running old firmware with Universal Plug n Play (UPnP) enabled; odds were that many of these were “NYP’d” (not yet patched). This was not an uncommon situation. On any given day, there are millions of systems on the internet, which would respond to a network scan with the port shown as open. These types of systems are perfect targets to become zombies for hackers to leverage and amplify an attack.

The three-part handshake

Applications that require reliable communications often leverage protocols built on the Transmission Control Protocol, or TCP. To achieve high levels of reliability, TCP uses a variety of control flags to communicate the state of a connection and validate receipt of data sent across a connection. When analyzed, these flags can provide valuable insight into network behavior, and be used to understand the nature of any malicious communications.

A core component of TCP is the “three-part handshake,” a mechanism used to validate that both hosts involved in a communication are aware and ready for data transmission. To initiate a connection, the requestor, or client, sends a TCP/IP packet to the requested host, or server, containing a single flag. This SYN flag asks the server to SYNchronize with the client. If the server agrees it will reply with a two-flag packet, SYN-ACK, as a way of both ACKnowledging the synchronization request as well as asking the client to synchronize with the server. If the client is still willing to participate in the communication it replies with a final single-flag packet, ACK, to let the server know the client is ready for further communications. Once the full handshake (SYN, SYN-ACK, ACK) has been completed, the two hosts may transmit data bi-directionally for as long as timeouts allow.

For a security analyst the TCP handshake can be a valuable way to quickly verify the state and legitimacy of network communication. When dealing with infected networks, we must frequently filter through large amounts of network traffic to identify potentially suspicious or malicious communications. One way in which connections can be eliminated, thus reducing the data corpus to review, is by searching for communications, which show only a full three-part handshake.

TCP sessions identified without these exact packets have a very low probability of data transfer, and therefore are unlikely to contain malicious activity or exfiltrated data. This is especially the case with DDoS-related attacks. Abnormalities in the handshake pattern or missing final ACK packets can be strong indicators of forged, or spoofed, traffic. In the case of highly segregated environments, even unsuccessful connections may be suspect and these flags can be used to triage types of traffic for later review.

Mitigating the UDP flood proved more difficult as it was destined for a port that this customer relied on for normal application traffic. Denying traffic to the UDP ports with a blanket statement would have also denied the legitimate user base. The threat actors knew exactly where to focus their attacks. Mitigation for the UDP flood had to be handled by an appliance, which would scrub the traffic in line, and subsequently drop packets that were not defined within the rule set parameters. A custom mitigation rule was created to match the payload signature, packet size, and destined port.

The SYN flood was also handled by a mitigation appliance, but would instead challenge incoming TCP connections. Spoofed source IP addresses wouldn't respond to the challenge and would be dropped. Legitimate user connections would reply successfully and make a full TCP connection. This particular mitigation strategy is effective but can cause collateral damage since there is no way of proving a user is legitimate without going through the same challenge mechanism in order to authenticate.

During the investigation, the RISK Team identified a known hacking group that was using the DDoS as a way to advertise their services. The threat actors stated that for a nominal Bitcoin fee, they could bring down any other application for an extended timeframe.

With the full mitigation stack in place, the DDoS attack's effectiveness subsided and services were restored eventually. As a result of the attack – and learning several hard lessons – my company was ultimately able to improve its overall security posture. Large-scale DDoS attacks can't fully be prevented, but having the right resources to battle them can drastically reduce downtime and hasten recovery.

Recommendations to add permanent Access Control Lists (ACLs) for incoming TCP connections were put in place. These entries followed a Request for Comment (RFC) standard for TCP flag combinations, and would drop invalid flag sets immediately without making it through the GRE tunnel and ultimately hitting the backend, which lies further downstream. The same recommendations were made for the UDP traffic. Unexpected ports would be dropped upstream and only legitimate destination port ranges would ever be allowed. Although this may seem like standard practice, networks change constantly and sometimes drastically and therefore should always undergo rule revisions.

In addition to adding traffic rules for inbound connections, the frequency of service validations and mock incident tabletop exercises were increased to every quarter. Having the capability to run a standard attack scenario every three months, without the same pressures of an actual attack, was now part of the standard regimen for all teams. These exercises allowed kinks to be worked out in a controlled environment.

Lessons learned

As the number of DDoS tools, IoT devices, and misconfigured systems increase, a security regimen that considers large-scale attacks is paramount. Having a strong security posture and remediation plan can drastically reduce downtime and hasten an organization's ability to respond and recover. Organizations would do well to consider the following baseline plan of action:

Mitigation

- Automate prefix routing to the DDoS provider and test the functionality periodically.
- Funnel advertised routes as intended.
- Increase bandwidth to essential networks.
- Use well-defined ACLs and firewall rules.
- Limit (half-open) connection rates.

Response

- Validate services to rule out unexpected complications during an attack.
- Conduct post-incident investigations.
- Conduct social media awareness campaigns.
- Document processes for handling DoS attacks.

It is vital to take a proactive approach to defending your network especially when your customers are using it. In these circumstances, additional security enhancements, like those listed in this exercise, can significantly reduce downtime for these types of attacks. Although it is a best practice not to engage an attack group, it is always advisable to keep an eye on social media feeds. Threat actors may brag about taking a company down or hint at attempting to do so. Any potential precursors to an impending attack will certainly reduce the element of surprise.