

2

STATE-SPONSORED FINANCIAL ATTACKS



As long as banks have existed, people have been trying to rob them. But until recently, criminals had to physically enter the bank, usually masked and armed, and use the threat of violence to demand money. Today, that is no longer the case. Over the past 10 years, the world has seen many high-dollar bank compromises in which the robber never stepped foot on the premises.

Computing technology and the internet have allowed banking to move from a brick-and-mortar access model, one that required customers to come to the bank to access their funds, to a system made of bits and bytes. In present times, we can remotely conduct banking from any internet-connected device. In fact, banking is more secure than ever thanks to this technology.

Unfortunately, connectivity has also provided criminals with new opportunities for theft. Banks today risk losing more money from a single

criminal operation than ever before. That's because a brick-and-mortar bank's financial loss is limited to the funds on hand at the branch. Online banking allows financial institutions to grant customers access to funds beyond those available at one physical location. While this enables banks to provide their customers with better service, it also means online attackers can steal vast sums of money.

Typical cybercriminals often don't have the means—or the time—required to execute attacks against financial institutions. Yet nation-state attackers pose a significant threat to financial institutions, as they have the resources and technological fluency to defeat robust cyber defenses. And remember, a government will have different motivations than a criminal. Here's something you may not have realized: financial gain isn't always the objective of these nation-state attacks. Prior to 2013, nation-state attacks against banks primarily caused denials of service. The governments that executed these operations—primarily Iran and North Korea—did so to make a statement, retaliate, or weaken the economic strength of the nation in which the bank operates. In 2013, after years of denial-of-service (DoS) attacks, nations began financial theft operations, as restrictions against these poorer nations were inhibiting their economies, motivating them to steal.

While it is now common to read about nation-state cyberattacks resulting in substantial economic losses, these attacks are still a relatively new threat. Understanding the evolution of these attacks helps explain how these nation-states became the financial attackers that they are today. In this chapter, we will discuss attacks against the financial industry and attackers' motivations and methodologies.

Distributed DoS Attacks Against Financial Institutions

On July 4, 2009, banking websites in the United States and South Korea became suddenly unresponsive; a massive cyberattack had infected a total of 50,000 computers, most located in South Korea, according to reports. The attack had used malware later named Dozer, which spread via phishing emails.

Perhaps surprisingly, the attacks did not attempt to steal money from the institutions. Instead, they crippled banks' functional capabilities by leaving them unable to provide services. Denying financial resources and services to consumers, it turned out, is an effective form of cyberattack: lack of access can often be as effective as outright theft. After all, banks usually protect and insure customer funds, but none of that matters during a bank outage. In these instances, consumers cannot use debit cards, withdraw money from automatic teller machines (ATMs), or even go to a branch to make a withdrawal. If you've ever gone to an ATM to withdraw funds and found that it was out of order, or attempted to use your debit card and had the transaction denied, imagine if that same problem prevented you from accessing your money for a week. It would likely make you think twice about how you handle your banking needs.

Hackers understood this, and the 2009 incidents were the first in a series of attacks designed to place doubt in the minds of consumers.¹ If enough people lost trust in banks and the financial systems behind them, the nation's economy could become affected. In the worst-case scenario, if consumers did not trust banks, they might begin to withdraw funds while ceasing to deposit money, causing a domino effect and potentially weakening a nation's economy. This is not as likely in countries with large, strong economies.

The Dozer Attack

The Dozer malware incident represents the first publicly known attack in which a nation-state targeted financial institutions, and it is widely attributed to North Korea. The phishing emails used in the attack contained an attachment that dropped several malware components onto victims' systems. From there, the attackers could leverage these compromised resources directly. These components included the following²:

- **W32.Dozer:** The mechanism that dropped the other malicious components.
- **Trojan.Dozer:** A component that provided the DDoS and *backdoor*, or remote access, functionality.
- **W32.Mydoom.A@mm:** A worm used for spreading the malware to additional victims.
- **W32.Mytob!gen:** A component that infected victims' systems, accessed their email contacts, and sent Trojan.Dozer to every entry in their address books. As this process continued, the rate of infection grew rapidly. This increased the number of resources involved in the DDoS component of the attack.

The attack involved other resources, too, such as botnets that attackers purchased or obtained through unreported means. Using other people's tools limited the chance of outsiders identifying their custom malware in the wild prior to the attack. The process of infecting thousands of systems would have provided defenders with an opportunity to discover and attribute the activity before the denial-of-service attack, lowering the chances of success. On the other hand, the attacker could purchase a botnet from cybercriminals with almost no risk of exposure.

The attack itself was clever primarily because it propagated itself using a worm that spread to other systems automatically. Once far more prevalent, this form of malware often appeared in the lower-level attacks of the mid-to-late 1990s and early 2000s. Even a simple worm could quickly share malware and other malicious components, leading to maximum infection with minimal overhead. Furthermore, attackers did not need to interact with any part of the systems manually.

Attackers conducted three waves of DDoS attacks between July 4 and July 9, each targeting a different set of websites, including the following finance-related domains: *banking.nonghyup.com*, *ezbank.shinhan.com*, *ebank.keb.co.kr*, *www.nyse.com*, *www.nasdaq.com*, *finance.yahoo.com*, *www.usbank.com*,

and *www.ustreas.gov*. While the attackers did not target financial institutions alone, this was one of the first instances in which a nation-state used cyber weapons to cause harm to the financial sector.

Unlike attacks originating from non-nation-state cybercriminals, the malware had unique characteristics: although it became active on July 4, the attackers configured it to terminate on July 10, ceasing the DDoS and launching the attack's second component.

Unfortunately for the victims, once July 10 arrived, the malware's final destructive act began. It began wiping data with specific file extensions from the systems. Then it erased their master boot records (MBRs), rendering the systems useless. Once done, the malware presented the message "Memory of the Independence Day"—a thank-you note of sorts from the attackers. This anti-U.S. message proved to be yet another clue that the attack did not originate from a cybercriminal.

At the time of the attacks, public speculation placed North Korea as the prime culprit. The attacks came as North Korea was conducting ballistic missile tests, despite previous sanctions against such tests. In 2014, the U.S. government confirmed the attribution.³

Ten Days of Rain

The next major DDoS attack targeting financial institutions occurred two years later. In its tactics and malware, the attack had many similarities to the 2009 attacks. More significantly, however, the 2011 attack replicated the three-phased operation of the Dozer incident. Later, other nation-states would adopt this attack model to use in their operations. Table 2-1 walks through this attack model.

Table 2-1: Nation-State Three-Phase Denial-of-Service Attack Model

Phase name	Attack details
Phase 1, "Bot infection"	In the first phase, the attackers infected hosts with malware, which built and powered the bot necessary for the DDoS phase of the attack.
Phase 2, "DDoS attacks"	The second phase used the system resources to target specific sites affiliated with organizations with a DDoS attack.
Phase 3, "Sabotage and destruction"	The third phase caused chaos, destroying systems and data, rendering them useless. The attackers also used this phase as an opportunity to display images and messages to the victim.

Once again, the public blamed North Korea at the time of the incident. In the years since, the U.S. government has discovered binary similarities in the malware used in the attack and other malware attributed to North Korea, bolstering this claim.

One of the differences between the 2009 and 2011 attacks is how the later malware, Trojan.Koredos, handled its configuration and DDoS target

data. The earlier Dozer malware communicated with a command-and-control infrastructure to obtain instructions and configuration parameters, such as the list of targets. This communication had to traverse networks between the victim and the adversary's infrastructure. By contrast, the Trojan.Koredos malware used in 2011 already contained the target list and attack parameters, making this external communication unnecessary. Automated defenses can identify when malicious activity is taking place on their network by identifying the network communications that originate from the malware itself. As the malware didn't require external communication, defenders had one less opportunity to identify and mitigate the attack.

Also, predetermined start and stop times were built into the malware itself. The attackers wanted the DDoS operation to last for 10 days. For this reason, the March 2011 attacks were dubbed the *Ten Days of Rain*.

During the attack, media outlets reported that some South Korean banks' servers crashed, and websites became unresponsive. According to the *Washington Post*, "30 million customers of the Nonghyup agricultural bank were unable to use ATMs or online services for several days." They stated that key data was destroyed.⁴

IRGC Targets U.S. Banks (2011–2013)

In late 2011, banks began to see spikes in the traffic affecting the performance of their systems and services, suggesting they had become the target of attackers. This initial activity likely constituted the attackers' dry run: a fire drill of sorts, used to test their ability to disrupt regular operations and discover if they could maintain an attack from one week to the next. But by September 2012, the activity had dramatically morphed from an engagement targeting a small subset of institutions to a major attack against many banks throughout the United States. The attackers had designed and organized their efforts to take down bank websites and resources concurrently.⁵

Once again, the attackers targeted the banks not for financial gain but to demonstrate their power. The DDoS campaign would continue through 2013, affecting approximately 50 U.S. financial institutions in one of the most comprehensive and lengthy DDoS campaigns known to date. Victims included well-known banks, such as JPMorgan Chase, Wells Fargo, and Bank of America.⁶

A Middle Eastern hacktivist group, the Izz ad-Din al-Qassam Cyber Fighters, soon took credit for the attacks. The group posted messages on Pastebin, like the one in Figure 2-1, that called for others to support its cause against the United States.⁷

We, Cyber fighters of Izz ad-din Al qassam will attack the Bank of America and New York Stock Exchange for the first step. These Targets are properties of American-Zionist Capitalists. This attack will be started today at 2 pm. GMT. This attack will continue till the Erasing of that nasty movie. Beware this attack can vary in type.
Down with modern infidels.

Figure 2-1: Izz ad-Din al-Qassam Cyber Fighters' message posted to Pastebin

Yet according to media reports at the time, sources involved with U.S. intelligence attributed the attacks to the Iranian government.⁸ This attribution relied on circumstantial evidence, and the reports did not name their sources, but the argument still held weight; DDoS attacks are common and do not require a significant degree of technical skill, which is why they are popular with hacktivist groups. However, no attacker had yet succeeded in sustaining such a lengthy, ongoing attack of this size against nearly 50 institutions. Post-compromise reports described how banks were hit with as much as 140 Gbps of data per second, making it the most powerful DDoS attack on record at the time. Moreover, the incident, which stretched over a year, proved longer lasting than any previously reported attack. That a hacktivist group would have been able to conduct and maintain a DDoS campaign of this scope is highly unlikely. Its magnitude suggests a state like Iran was behind it.

If this was not the work of a hacktivist group but the nation-state of Iran, then the Izz ad-Din al-Qassam Cyber Fighters attribution functioned as a coordinated disinformation campaign. Although not the first time a nation used disinformation to provide plausible deniability, it is one of the most public instances coming from Iran.

In March 2016, the U.S. government issued a federal indictment against two organizations, ITSEC Team and Mersad Co. The indictment described these as “private computer security companies based in the Islamic Republic of Iran” that performed “work on the behalf of the Iranian Government, including the Islamic Revolutionary Guard Corps (IRGC).” The affidavit charges the organizations—and specifically seven Iranian citizens—with infecting computers, building a botnet, and conducting a DDoS campaign against financial institutions from 2011 through 2013. Figure 2-2 is the image released by the FBI of the individuals charged in the attacks.⁹

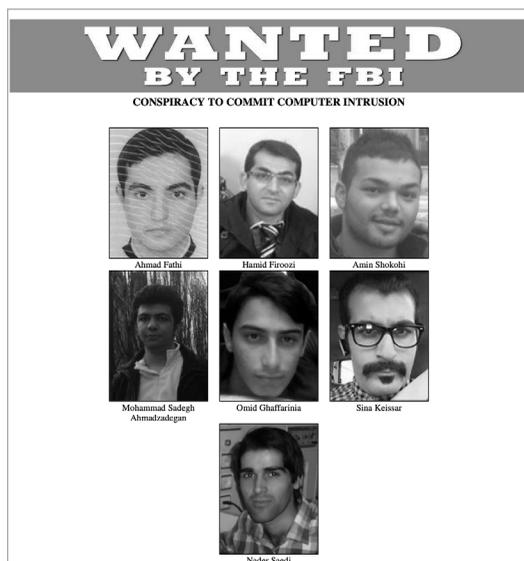


Figure 2-2: Individuals wanted by the FBI for taking part in Iran-based DDoS attacks against financial institutions

These, allegedly, are the faces of the DDoS attack. Still, many others, including higher-ranking individuals associated with the IRGC, likely took part in the attacks or at least had relevant knowledge of it. The U.S. government probably released the indictment publicly to send a message to the Iranian government, as it is unlikely that the United States will ever apprehend these men. The United States has no jurisdiction in Iran, nor will the Iranian government cooperate in convicting operators it hired to support its operations.

Public and media speculation has proposed that the attacks came in response to sanctions against Iran's nuclear program, as well as retaliation for the Stuxnet attacks against Iran's nuclear facilities in 2010.

DarkSeoul

On March 19, 2013, the cybersecurity firm Trend Micro detected a wave of spear-phishing emails targeting South Korean financial institutions.¹⁰ The following day, banks and media organizations began reporting widespread outages; the malware had destroyed their infrastructure, rendering their systems and resources useless. Trend Micro released a report stating that the “websites of several banks may have been compromised and exploits [were] used to plant backdoors on the systems of [website] visitors.” Avast, an antivirus vendor, published its own blog documenting what it believed was a strategic web compromise geared at South Korean banks.¹¹

At the time, neither vendor had all of the details of the attack correct, as both had come across something much bigger than they originally realized. Cybersecurity officials blamed China at first: the attack relied on adversary infrastructure located in China, and Chinese names were found in the malware. Future evidence would later prove these attributions incorrect, serving as an excellent example as to why it is smart to use more reliable supporting evidence before making public attribution assessments.

In addition to misattribution muddying the waters, the attackers took steps to misdirect blame through diversion and misinformation. They created two social-media-based hacktivist groups, the NewRomanic Cyber Army Team and Whois Team. These groups claimed responsibility for the attacks by posting messages, such as the one in Figure 2-3, on defaced websites and victim computers.¹²



Figure 2-3: “Whois Team” message taking credit for 2013 DDoS attacks

Nobody had heard of either group prior to the 2013 attacks, leading many to believe, correctly, that someone had simply invented the personas; the groups produced no posts or affiliated social media accounts before or after March 2013. In fact, South Korean government officials have since claimed the attack originated from North Korea. Unique code found only in North Korean malware supported their attribution.¹³

Like previous North Korean DDoS attacks, attackers used spear-phishing emails, as well as compromised websites, to infect victims' systems with first-stage malware. This attack also used a destructive wiper malware, although this time, attackers did not begin by leveraging a botnet to take down websites and servers. Instead, the attackers directly infected the intended target with the wiper, which itself functioned as a denial of service. By destroying systems and data in the targeted organizations, such as South Korean financial institutions, the attack made critical services unavailable and therefore had the same effect as the previous DDoS campaigns.

In addition to the infection vectors mentioned, attackers used a third, more creative, and especially effective vector: a software update mechanism. This allowed them to bypass target defenses and stealthily plant malware onto many systems within the targets' infrastructure. The attackers knew that South Korean financial organizations would likely use South Korean security vendors to protect their assets. As it so happened, at least one financial target in this attack used software from AhnLabs, a South Korean vendor, for both its antivirus and patch-management solutions. Thus, before deploying the wiper malware on victims' systems, the attackers gained control of an account with administrative access to AhnLabs' patch management software within the targets' local environment.¹⁴

AhnLabs itself was never compromised in these attacks, as the first reports indicated. Instead, attackers obtained the AhnLabs credentials from the victims' local environment. Attackers can do a lot of damage when they obtain administrative privileges in an environment with many unpatched systems; the vendor's patch management software provided updates to almost every system within the client environment. The attackers used this to bypass the targets' firewalls and security defenses, delivering malware instead of software fixes. By disguising it as a software update, attackers silently distributed the wiper throughout the targets' infrastructure, where the infection spread to other targets through a variety of means.¹⁵

At 2 PM local time, the wiper executed across seven victim organizations: four financial institutions and three media companies. Attackers had designed it to destroy the master boot record on the targeted systems, preventing the systems from starting up. The 2009 and 2011 DDoS attacks attributed to North Korea had also done this. In those attacks, however, the malware had simply deleted the boot record, and while not easy, it's possible to recover from such a deletion. The wiper malware in the 2013 attacks took an additional step of overwriting the record, and all data on the associated drive, with the strings "PRINCIPES" or "HASTATI." By overwriting the data instead of just erasing it, the attackers made it much harder, if not impossible, to recover the lost contents. Once the malware finished wiping and overwriting, the malware forced the system to reboot, rendering it useless, since the malware had removed

all data that the system needed in order to boot. The attack affected at least 48,700 systems upon reaching the predetermined end time.¹⁶

Once again, bank customers found themselves unable to withdraw or deposit funds through ATMs. Additionally, employees were unable to use bank terminals to assist customers, leaving many customers with no access to their accounts. Bank websites experienced intermittent outages, or were slow to respond, and the affected broadcast companies reported taking entire networks offline.

Two other technical details make this attack unique. The first is that the attackers tailored the malware to infect various operating systems. Corporate environments most commonly use Microsoft Windows as their operating system of choice; Unix-based systems, on the other hand, are prominent in backend banking platforms. The wiper was capable of erasing Windows systems as well as Unix-based ones, such as AIX, HP Unix, Linux, and Solaris, which often authorize and coordinate information exchanges within banking transactions.

The second unique aspect of these attacks is that the attackers tailored the malware to look for and disable specific antivirus programs running within the target victims' environment. If the victims had installed either Hauri or Ahnlabs antivirus software on their systems, the wiper component activated itself only after disabling the security software, ensuring its successful execution.

The sophistication of this attack is worth emphasizing. The malware included several nation-state attribution hints in its design to throw off security researchers. It targeted multiple operating systems and relied on various delivery vectors, antivirus evasion, and mitigations, showing the attackers put time, effort, and resources into the attacks prior to the campaign execution. Finally, using fake personas to take credit for the attack is a tactic that cybercriminals or hacktivists rarely use. All of these elements of the campaign are hallmarks of a nation-state attack.

Russian Attacks Against Ukraine

Although we won't discuss the topic in detail in this chapter, Russia has conducted similar attacks to those discussed thus far, resulting in a DoS of banks in Ukraine.¹⁷ For example, in 2014, Cyber Berkut, a nation-state group with strong ties to Russian intelligence, forced PrivatBank—the largest Ukrainian commercial bank—to shut down operations.¹⁸ The attackers compromised the bank and then released both sensitive customer and bank operational data to several public websites, including Twitter and V.K., a Russian social media platform. The data included customer names, addresses, and account balances, as well as engineering and infrastructure information specific to the bank's internal network. The final nail in the coffin for the bank came when the attackers instructed bank customers to remove their money from the bank or permanently lose access to their funds. The bank never truly recovered from the attack or the resulting loss of customers, who likely lost faith in the institution's ability to protect their money.

Within two years, the disaster forced the Ukraine government to take over the bank's operations, preventing bankruptcy and removing its commercial

interests, to make it a 100 percent state-owned institution.¹⁹ The cyberattack against the bank may not have been the only cause, but it contributed to a ripple effect across the nation's economy, forcing the bank's nationalization.

Billion-Dollar Robberies

The world didn't begin to see large-scale financial thefts until 2013, when cyberattackers, likely from North Korea, stole funds from Sonali Bank in Bangladesh. The first confirmed North Korean financial theft would take place in 2015, though many similarities exist among the tactics and behaviors present in the 2013 and 2015 attacks.

These thefts likely came as a consequence of economic sanctions imposed on North Korea. The sanctions, which aimed to prevent the growth of the state's military and nuclear capabilities, kept North Korea from trading with other countries, including importing critical oil and gas, therefore forcing North Korea to rely on homegrown assets and resources.²⁰ To remain relevant on the world stage—and not starve—North Korea has had to look for more creative ways to grow its economy.

Unfortunately for the financial industry, one of North Korea's primary responses to the sanctions has been cyberattacks. Its cyber campaigns have successfully stolen hundreds of millions of dollars.

SWIFT Attacks

Many of these financial thefts began with the compromise of the Society for Worldwide Interbank Financial Telecommunication (SWIFT) messaging system. SWIFT is software that financial organizations use to communicate transaction information with each other.²¹ North Korea obtained access to the organizations' internal SWIFT systems, as in the following attacks, which cybersecurity officials have either attributed to North Korea or matched with tactics present in known North Korean attacks (see Table 2-2):

Table 2-2: Timeline of Financial Institutions Targeted by North Korea

Year	Country	Institution
2013	Bangladesh	Sonali Bank
2015	Ecuador	Banco del Austro
2015	Vietnam	Tien Phong Bank
2016	Bangladesh	Bank of Bangladesh
2017	Nepal	NIC Asia Bank
2017	Taiwan	Far Eastern International Bank
2018	Mexico	Central Bank and Banorte
2018	India	City Union Bank
2018	Chile	Banco de Chile

In some instances, the attackers used spear-phishing emails to distribute malware; in others, they used watering holes. But the attackers never compromised the SWIFT organization itself. Instead, they exploited vulnerabilities in the client-side systems at banks, which enabled attackers to alter systems utilizing SWIFT messaging transactions. This is important to mention, as a large number of financial organizations continue to rely on the integrity of SWIFT. Today, SWIFT itself remains trustworthy.

The North Korea Financial Theft Model

On June 8, 2018, the U.S. Department of Justice issued a criminal complaint against a North Korean citizen named Park Jin Hyok. The complaint documented several computer-related crimes, including hacking, that Park conducted along with unnamed individuals. The complaint provides an inside look at the hacking operations of North Korea, one of the most notorious nation-state attackers to date. It also offers extremely useful details for a defensive perspective. This section draws on this information.

The staged attack model listed here originates from details within the Department of Justice's criminal complaint, in conjunction with research and publicly available analyses from security vendors.²² It involves the following phases: reconnaissance, initial compromise, observation and learning, enumeration and privilege escalation, preparation of the staged environment (account and resource creation), execution of fraudulent transactions, and deletion of evidence.

While some of the malware and tactic details varied from one attack to another, North Korea continued to use the same phased attack described earlier. It's fair to conclude that North Korea will use the same approach for as long as it succeeds.

Reconnaissance

The attackers spent considerable time performing reconnaissance. For example, Park conducted online reconnaissance "a year before the cyber-heist at Bangladesh Bank."²³ During this stage, the attackers would gather information about the bank's public-facing infrastructure, as well as associated email addresses. Park researched the target bank's website and employees, including their social media accounts. In some instances, the attackers used services that specialized in "locating email accounts associated with specific domains and companies."²⁴

Attackers collected email addresses to create target lists for use in the next phase of the attack. In some instances, the attackers created spoofed accounts that mimicked someone known to the target. In others, the attacker created email addresses to register social media accounts. Attackers leveraged these social media accounts in later stages of the attack. Furthermore, attackers also mapped out the target's public infrastructure, likely in an attempt to identify any vulnerabilities that they could exploit to gain access to the victim's environment in later stages as well. Park also researched specific vulnerabilities to identify how to exploit them. Presumably, these

were vulnerabilities he identified when conducting reconnaissance into the Bangladesh Bank's infrastructure.

In addition to these factors, attackers created and staged accounts and online personas during the reconnaissance phase of the attack. They created email accounts from free, publicly available webmail platforms such as Gmail. Later in the process of the attack, these accounts interacted with bank employees and sent spear-phishing emails.

Initial Compromise

Multiple North Korean financial theft campaigns used social engineering in the form of spear-phishing emails to compromise and gain access to the target's environment. Attackers tailored these spear-phishing emails to target the individuals and accounts that they had identified during reconnaissance. According to U.S. federal investigators, North Korean hackers crafted emails in several high-profile bank attacks that were “highly targeted, (and) reflect the known affiliations or interests of the intended victims, and are crafted—with the use of appropriate formatting, imagery, and nomenclature—to mimic legitimate emails that the recipient might expect to receive.”²⁵ In other words, the attackers spent time and resources to make the email specific, relevant, and appear legitimate to the targets.

Once compromised, attackers used the email accounts to send spear-phishing emails to other bank officials from legitimate accounts. This aspect of familiarity added legitimacy to the emails. The attackers often were not interested in compromising additional recipients; however, they included them, so the actual target saw familiar email addresses in the “To” or “CC” line of the email. This tactic demonstrates the level of detail and planning the attackers put into their spear-phishing emails.

Companies often use public-facing email addresses that are not attached to a specific individual. Instead, a group or an administrator at the organization monitors these public-facing email addresses. A typical example of this is when companies use a single email address to receive résumés and other types of correspondence. At Bangladesh Bank, the attackers recognized such an email address as an opportunity to submit a résumé weaponized with malware. Examples within the criminal complaint included links in the body of the email requesting that targets click to view a résumé. When the targets clicked the link, malware compromised their systems, providing attackers with access to both the system and the environment.

Other North Korean compromise attempts included the use of emails mimicking alerts or notifications from social media and service providers such as Google and Facebook. For example, attackers utilized standard emails alerting users when someone accessed their account from a new location. The fraudulent emails mirrored legitimate ones by including the same text and images. The primary differences between the two were the sender address—which attackers also often spoofed—and the URLs within the email. Attackers made sure to obfuscate the links in order to appear legitimate, but these links took the victims to attacker-controlled infrastructure to infect them with malware.

Financial institutions suffered from attacks other than spear-phishing campaigns, however. In 2016 and 2017, legitimate financial-themed websites that other banking companies and individuals often visited succumbed to infection. These websites then infected site visitors with custom malware. For example, attackers compromised the website of the Polish Financial Supervision Authority, and the website later infected financial organizations in Poland.²⁶ The attackers knew that many other banks in this region would often visit the website. Similar attacks occurred around the same time, affecting the site of a Mexican financial regulator and a bank in South America. Each attack compromised systems and resulted in the website serving malware to website visitors. Later, analysis of the malware distributed by the compromised sites showed an overlap in code only previously seen in North Korean malware.

Observation and Learning

In all of the North Korean–attributed financial attacks, the attackers spent time learning the local environment. Based on the behaviors seen across multiple intrusions, North Korea is a patient attacker that spends considerable amounts of time within the targets’ environment before executing the financial theft phase of the attack. In some cases, the attackers spent several months observing and learning the systems and how they connect and interact with other banking resources.

For example, a unique attribute of these attacks is the amount of time the North Korean attackers spent learning the banks’ policies and procedures. Here, the objective for the attackers was to better understand how employees handle and conduct financial transactions. This is notable because, except for nation-state espionage campaigns that were not a major concern to financial institutions at the time, it was generally unheard of for an attacker to spend time learning the targets’ employee policies and procedures. Doing so, however, is another example of the planning and patience the attackers put into these operations. This also illustrates the differences between a typical financial attacker and a nation-state attacker.

North Korea’s diligence in learning the banks’ noncyber policies paid off. Two of the targeted banks, Tien Phong Bank (Vietnam) and Bank of Bangladesh, archived SWIFT transactions differently than most financial institutions. Bangladesh Bank printed paper copies of SWIFT messages. Hard copies of the transactions provided a physical record archived at the bank. Tien Phong Bank, however, stored electronic PDF versions of the messages on a third-party server. It used FoxIt Reader, an application for managing digital documents such as PDFs, to convert SWIFT message details into PDF records. The attackers identified this process and developed malware that would infect the bank’s systems when bank employees attempted to access the PDF software by replacing that application with a weaponized version of the software.

If the attackers had tried to implement this at Bangladesh Bank, it would not have worked. This is because the bank used printed copies to archive transaction messages. Alternatively, at the Vietnamese bank, if the

attackers had attempted to print hard copies instead of saving the messages as PDFs, it would likely draw attention to their activities. Taking the time to learn each bank's unique business processes allowed the attackers to identify creative ways to further infect and quietly execute fraudulent transactions. More importantly, the attackers used the information to blend in with legitimate bank activity.

Enumeration and Privilege Escalation

The attackers also used various hack tools (often publicly available) to enumerate the victims' environments. The goal of enumeration was to identify computers the bank used to send and receive messages via the SWIFT communication system.²⁷ As part of their security practices, the targeted institutions implemented a *segregation of duties* policy within their environments. This is a practice that prevents any one person from having complete access to critical business systems and functions within the environment. Unfortunately, this did not prevent the attackers from gaining the necessary access to attempt fraudulent financial transactions. It did, however, increase the difficulty of the attack. The attackers needed access to multiple protected accounts to get into various systems and segregated networks before infiltrating the accounts and systems associated with SWIFT transactions.

Many of these administrative accounts fell into attackers' control via using credential-collecting hack tools, such as keyloggers, or through spear-phishing emails sent from legitimate internal bank accounts. One such keylogger present in the Bangladesh Bank heist hid within the `C:\Windows\Web\Wallpaper\Windows\` directory on a compromised host, indicating the malware may have been delivered through an attachment mimicking desktop wallpaper.²⁸

Preparing the Stage

To continue operations and stage the target environment, the attackers needed to maintain an undetected presence. The malware's communication traffic could have caught the attention of defenders as it actively communicated with both internal victim infrastructure and adversary command-and-control servers.

In an effort to hide their activity, attackers used what has been described as a "custom binary protocol designed to look like 'TLS' traffic" to encrypt the malware's communications.²⁹ TLS, short for Transport Layer Security, is an encryption-layer protocol that protects network communication traffic such that it cannot appear as clear text while in transit. The attackers used a version of the TLS protocol that had a fake TLS header. The TLS header leveraged a unique cipher suite with a hardcoded array, altering network traffic at the encryption level, making it difficult to detect. Then the attacker created a second version, which also used a fake header; however, instead of a hardcoded array, the cipher suite used a random cipher. These were then appended to the command-and-control communication traffic generated by the malware. A *cipher suite* is comprised of algorithms

used for cryptographic operations, such as encryption and decryption, and allows for key exchange and other authentication procedures that banks commonly use today to secure traffic between communicating hosts.

The attackers built the encryption protocol into a custom-developed backdoor known as NESTEGG. Without the proper encryption key or an understanding of the custom protocol, nobody could decrypt traffic originating from the infected system. Since the communication traffic appeared similar to legitimate TLS traffic, the attackers were able to communicate with command-and-control infrastructure covertly.

The attackers added another level of complexity by having the NESTEGG backdoor run in memory on the victim system. We call malicious code that runs exclusively in memory on the victim's system *fileless malware*. The benefit of this design is the malware can go undetected, since it's not written to, or present on, a physical drive; it executes and runs commands directly in memory. Most security products monitor and detect files as they write to the hard disk of the protected system.

The drawback of fileless malware is its lack of persistence. Since the disk is not written to, fileless malware can be deleted if the infected system reboots or restarts. The NESTEGG malware, however, addresses this shortcoming by monitoring the victim system to detect shutdown and reboot functions. When it identifies either of these events, the malware installs a copy of itself onto the victim's hard drive to reinstate itself once the operating system restores. After rebooting and reinstalling, the malware deletes the copy written to the hard disk and once again exists only in memory on the victim system.

NESTEGG had various other notable functions, such as “acting as a proxy to send commands to other infected systems, and [accept] commands to upload and download files, list and delete files, start, and terminate processes.”³⁰ These capabilities allowed the attackers to stage, prepare, and further compromise the banks' systems and networks. Specifically, the attackers placed malware on various systems involved with processing the banks' financial transactions.

Execution of Fraudulent Transactions

Up to this point, the attackers had gained access; observed bank systems, applications, and processes; and staged malware throughout the bank's network. Using the malware and information gained, the attackers were able to acquire various types of administrative accounts. Typically, no single entity would (or should) have complete access to the systems and components used to conduct a bank's financial transactions. However, these attackers used vast resources generally not available to typical criminals to obtain all the credentials necessary to authorize financial transactions.

Next, the attackers used the accounts to log into the SWIFT Alliance application, a message interface application, to conduct financial transactions. The SWIFT systems are usually separate from other bank networks, and network segregation, enforced with routers and firewalls, protects the systems. In the Bangladesh Bank heist, however, the bank's infrastructure

did not meet the security standards that should have been in place. In a report titled “North Korean Cyber Capabilities,” the U.S. Congressional Research Service noted the following³¹:

Bangladesh’s network may have been particularly vulnerable, as it reportedly lacked a firewall to protect against outside intrusion.

Of note, in some of the North Korean financial attacks, the attackers obtained access to legitimate accounts, while in others, they created new ones. This included the operator accounts necessary to access the local SWIFT Alliance application. The Alliance application is a “messaging interface (that) allows banks and market infrastructures to connect to SWIFT” and allows various financial institutions to create and confirm financial transactions.³² If the targeted institution had proper security controls in place, the creation of the operator accounts should have appeared to the institution as an uncommon or unusual event. In addition to this, the attackers unsuccessfully attempted to log into the Alliance application. Unfortunately, neither the creation of the operator accounts nor the failed login attempts alerted anyone, and the attackers gained complete access to the bank’s local SWIFT systems.

As previously mentioned, the attackers likely selected banks in countries or regions they believed to have weaker or less developed technology security standards. Between using printed physical copies of SWIFT transactions and not securing SWIFT systems, it is fair to say Bangladesh Bank was an easier target than many other financial institutions.

At this point, the attackers began to execute financial transactions. The transactions appeared legitimate, given that an account with valid access to the SWIFT system created and authorized them. From an outside perspective, as other banks involved in the transaction would view it, these were legitimate transactions made with the proper authorization and access. Before 2013, this type of attack had either not taken place or not been publicly acknowledged, so there was no reason to doubt the legitimacy of the transactions. In February 2016, the attacker-created SWIFT operator accounts attempted at least 35 transactions. In total, North Korea tried to steal nearly 1 billion dollars from the Bangladesh Bank.

Timing the Transaction Attempts

According to a 2019 public report that SWIFT published, the attackers documented the time of the fraudulent transfers at the Bangladesh Bank.³³ A pattern appeared: the transactions primarily occurred after working hours, between 11 PM and midnight in the local time. The report also documented the time of the attackers’ financial transactions at other banks believed to have been targeted by the same North Korean attackers. Almost every attack occurred between 9 PM and 4 AM local time, when the banks were closed.

The second pattern present in several of the bank attacks deals with the dates of the attacks. In several incidents, the attackers attempted fraudulent transactions on holidays, when banks were closed. By conducting the

transactions later in the evening to early morning and on holidays when bank employees are less likely to be present, the attackers had an increased chance of success.

Deleting Evidence and Covering Tracks

Methods and procedures varied for handling records associated with SWIFT transactions at targeted banks. From an attacker's perspective, if a bank employee or the bank's systems identified the transactions, this could give away their operation. To address this, the attackers designed features in their malware to delete files and other evidence left during the compromise. For example, a forensic investigation of compromised bank systems identified signs that the attackers had attempted to remove entries from system logs. Another common tactic seen across all the financial attacks was to delete malware from the infected systems once it had completed its given task. Specifically, multiple North Korean malware variants such as Contopée, NESTEGG, and SierraCharlie included a "secure delete function." However, the way the malware achieved this differed from one variant to another. Additionally, while not always successful, the attackers attempted to remove evidence of login attempts to the SWIFT Alliance application and its associated database(s).

It is highly likely the attacker behind the SWIFT banking attacks is the adversary behind the 2014 Sony Pictures Entertainment attacks. Components in the malware, such as the secure delete function and the custom cipher protocol, may have been initially designed for the Sony attack and then modified or updated for use in the bank attacks between 2015 and 2018.

Bank of Bangladesh Response

The Federal Reserve Bank of New York received the attacker-generated transaction requests. These transactions processed money transfers to accounts in the Philippines and Sri Lanka. Fortunately for the Bangladesh Bank, the total amount of the funds stolen was far less than the 1 billion dollars that attackers had requested. Ironically, these attackers, who spent a year carefully planning every detail of the heist, made a mistake in the most critical phase of their attack: they misspelled the name of a destination bank in one of the transaction requests. The attackers spelled "NGO, Shalika Foundation" as NGO Shalika "Fandation." This simple spelling error was enough for one of the banks routing the money to catch the activity.³⁴ When the routing bank identified the misspelling, it contacted Bangladesh Bank, which immediately terminated the transaction.

The North Korean attackers would have stolen almost a billion dollars, but according to media reports, the Federal Reserve had also contacted the Bangladesh Bank because of the unusually large amount of transfer requests and funds going to private organizations, such as the NGO. The bank stopped the pending transactions. In total, the banks managed to retain between \$850 and \$870 million by stopping these

transfers prior to reaching attacker-controlled accounts. Still, the attackers successfully made away with approximately 101 million dollars from Bangladesh Bank.

FASTCash: A Global ATM Robbery

On October 2, 2018, the Department of Homeland Security (DHS) released the US-CERT Technical Advisory, alerting financial organizations to a new attack that used custom malware known as *FASTCash*. According to the advisory, attackers had been working on this strike, which targeted financial organizations located in Asia and Africa, since at least 2016. Additionally, the U.S. government attributed the attack to Hidden Cobra (a name the U.S. government gave to North Korean nation-state attackers).³⁵ Following the alert, several security vendors produced research on the operation. One report found an overlap in the code FASTCash used and several other North Korean variants of malware, further supporting the attribution.³⁶

The Planning

North Korea is known for its creative and elaborate ways of stealing money to support its operations. This creativity came into play here, too, when a number of their bank heists only partially succeeded, as other routing banks flagged the financial transactions and stopped them while in transit. To get around this, the attackers developed a plan that would remove the routing banks from the process, eliminating the chance for them to claim that something was awry.

Many of the tactics seen in the previous North Korean bank attacks appeared in the FASTCash campaign. To gain access to the bank's environment, the attackers sent spear-phishing emails to bank employees, which infected their systems with custom malware. Once attackers obtained access, they spent time observing the victims' environment before attempting to steal funds. During this observation period, they also escalated their level of access and identified vulnerable areas of the bank's infrastructure.

For the FASTCash attacks, the attackers identified banks in Asia and Africa that used an outdated, unsupported version of AIX, a UNIX-based operating system that IBM created. Since FASTCash is not effective against current versions of AIX, it is unlikely that North Korea developed the malware before the breach. Instead, they took advantage of the opportunity once they discovered the vulnerability. Experts theorize that North Korea targets smaller banks in countries with weaker economies, as these are likely to have less operational funding and therefore are more likely to have outdated software and security controls.

The Execution

By exploiting the backend financial systems that banks used to process and authorize cash disbursements, the North Korean attackers were able to approve transactions that liquidated ATMs across 30 countries. The breadth of FASTCash left experts with little doubt that this was not the work of a

typical attacker but of a nation-state. While the malware's functionality varied, it shared similar design principles with the malware present in previous bank attacks. For example, both the SWIFT attacks and the FASTCash campaign used malware designed to interact with bank transaction authentication services; the earlier malware compromised the banks' SWIFT system to authorize the transfer of funds to attacker accounts, and FASTCash did the same with transactions involving ATMs.

Here, broadly, is how FASTCash works: when bank customers withdraw money from an ATM, they insert debit cards and enter their PINs. The ATM uses the PINs to authenticate the cards' owners. Once authentication is complete, the ATM reaches out to software called a *payment switch application*, or *switch*, to process the customer requests. The switch checks if there are sufficient funds in the account and then tells the ATM to either approve the transaction and dispense cash or deny the request.

The FASTCash malware prevented the switch from transmitting and processing fraudulent requests generated on the ATM. To do this, it monitored ATM transaction messages for account numbers the attackers had obtained in the preliminary phase of the attack. If the malware recognized the account number, it responded to the ATM with a transaction approval message, imitating the payment switch. The ATM believed the request to be genuine and thus dispensed cash without ever sending the request to the actual switch. In some reported instances, ATMs dispensed cash until they ran out, because the approved request exceeded the funds on hand.

Later, investigators would learn that the attackers had such a strong foothold in the targeted banks' networks that they had been able to create fraudulent bank accounts using legitimate systems. The attackers had given these accounts balances of zero to avoid drawing attention; as the malware acted as a middleman, preventing the actual switch from receiving the request, the accounts didn't need to be funded for the attack to work. Eventually, investigators matched these accounts to those within the malware that liquidated ATMs.

At least two times, once in 2017 and again in 2018, North Korea used FASTCash to execute coordinated simultaneous fraudulent transactions. In 2017, North Korea stole funds from multiple banks at the same time in more than 23 countries, in addition to the 30 countries targeted in 2018. One of the banks, located in Africa, came under attack in 2018 and could not return to normal operations for several months. Systems supporting ATM and point-of-sale services damaged in the attack left the bank unable to support their customers' business operations.

In 2020, bank heist operations continued and evolved. North Korean attackers had several years of successful attacks targeting bank payment switches with FASTCash malware. However, the adversary faced bank technology limitations. Banks use different systems to perform transactions. Not all banks used the vulnerable version of AIX, limiting the institutions North Korea could target. To expand the target base of banks in which they could attack, North Korea evolved and adapted, creating new versions of FASTCash designed to exploit Microsoft Windows servers in addition to AIX. As of September 2020, FASTCash operations attempted to steal more

than 2 billion dollars.³⁷ Additionally, the attacker began using wiper malware, destroying bank systems as a distraction while attempting fraudulent transactions. Between the expanded infection capabilities of the malware and additional destructive tactics, FASTCash operations have become one of the largest growing threats to financial institutions.

Odinaff: How Cybercriminals Learn from Nation-States

Earlier in this book, we pointed out differences between ordinary cybercriminals and nation-state attackers. Few cybercriminals are capable of the persistence, patience, and planning used in the engagements covered in this book so far. Unfortunately, there are always exceptions.

The North Korean SWIFT attacks made global headlines in 2016, garnering the attention of an organized cybercrime group named Odinaff. That year, security researchers revealed what they had discovered of the tactics, techniques, and procedures used in the SWIFT attacks to compromise the banks. This information has helped better defend against these incidents. But it also provided criminal attackers with a roadmap for future bank compromises.

Believed to originate from Eastern Europe, Odinaff successfully exploited banks with its own malware. It relied on tactics first seen in North Korean attacks, and current intelligence suggests that the group successfully stole millions of dollars from financial institutions.³⁸

As an initial attempt to gain access to the banks' systems, the attackers injected malware into a popular administrative tool called *AmmyAdmin*. They hoped bank administrators would download it, effectively infecting themselves. To do this, the attackers compromised the legitimate *AmmyAdmin* website—an attack that may sound elaborate, but in fact, criminals have frequently compromised the same site to distribute commodity malware.

NOTE

The website used to host AmmyAdmin has been known to distribute remote access trojans, exploit kits, and ransomware. Due to this risk, you should not visit the hosting website or download this tool.

While the *AmmyAdmin* tool might perhaps have functioned as an effective infection vector, the attackers likely realized it gave them no control over who downloaded the application. This risked infecting many unintended victims. It also exposed them to unwanted public attention. Probably for this reason, the attackers switched to the spear-phishing emails, which allowed them to choose their targets.

Odinaff's spear-phishing emails were nowhere near as sophisticated as North Korea's. Although targeted, the phishing campaign used a generic email template directing recipients to click a URL in the body of the email. The URL would then download a malicious payload. The attachment, however, did not infect victims if they opened it. Instead, victims had to open a compressed file that required the target to enter a password included in the email text. If victims followed the attackers' instructions, the archive would

decompress and present the target with a Microsoft Office document. Once victims attempted to open the document, the attachment presented them with the option to enable macros. If the target did not enable macros, the infection would fail.

Only if victims followed all of these steps did the first-stage malware, known as Trojan.Odinaff, compromise the system, providing the attackers with initial access to the victims' environment. That the attack required so many active steps on the part of the victims points to its precarity; if the targets had become suspicious of the emails, or perhaps the unusual requirements necessary to open the attachment, the attack would have failed. It may seem hard to fathom that anyone would fall for such a scheme. Yet it happened more than once, in attacks across several banks.

The Odinaff malware provided basic backdoor functionality, issued shell commands, and downloaded and executed additional malware. It used something called a mutex, hardcoded into the binary itself. A *mutex* is an object in the code used as an identifier. In this case, the identifier revealed whether a system was already infected. If it was, the malware halted execution. This prevented multiple infections on the same host from taking place, which would have tied up additional resources and potentially drawn unwanted attention. The malware also used a hardcoded proxy to connect to command-and-control servers, making it difficult for defenders to identify outgoing traffic.

Once in the victims' environment, the attackers would review the infected victims and identify systems of interest. They then used Odinaff's malware to download the stage-two malware, known as *Backdoor.Battle*, onto the subset of high-value systems of interest. (Researchers coined the name *Backdoor.Battle* after a string they found in the malware code containing the term "BATLE_SOURCE.")³⁹ The Battle malware ran malicious payloads in memory on the victims' systems, and it created a reverse shell, launched from a batch file, between it and the attackers' infrastructure.

The Backdoor.Battle malware was designed and developed using common penetration-testing software, such as the red-team tools *Metasploit* and *CobaltStrike*. The Metasploit framework identifies vulnerabilities and executes exploitation code against them. CobaltStrike functions with Metasploit to provide various post-exploitation and attack-management capabilities. Penetration testers commonly use both for legitimate security assessment exercises. Unfortunately, cyberattackers also use this tool to find and exploit weaknesses in victims' environments.

Odinaff's attack shared another tactic with those of nation-states: the use of tools already present in the victims' environment. Using legitimate administrative tools and applications already present on the system, the attacker can weaponize Microsoft Windows operating system binaries. This tactic, known as *Living Off the Land Binaries (LOLBins)*, allows attackers to hide malware in legitimate system binaries often whitelisted by security tools. When a binary is whitelisted, tools such as antivirus and endpoint detection software will not detect the file as malicious. Whitelisting prevents security tools from removing or quarantining the legitimate operating system resources that

could affect system functionality. Knowing this, attackers take advantage of the legitimate resource to use in attacks and avoid detection.

The Odinaff attackers used Windows administration software, such as PSEXEC, Netscan, and PowerShell. When the attackers needed to fulfill a capability unattainable by tools present in the victims' environment, they relied on publicly available hacktools instead of custom ones. A growing trend in cyberattacks, this strategy makes discovery and attribution more difficult. For example, both criminal and nation-state attackers have used the hacking tool Mimikatz against banks, because it is freely available, effective, a favorite of legitimate red teams, and impossible to attribute.

Using Batle, the attackers learned everything they could about the victims' environment. They spent time monitoring banks' activities and exploring the systems and infrastructure. Specifically, the Batle malware included the ability to capture keystrokes and images of users' screens in 5- to 30-second intervals. It then saved the output to a disk, where attackers could retrieve and study the captures. This allowed criminal attackers to learn the banks' processes and technical procedures for the execution of financial transactions. Another capability of the Batle malware—again, modeled after the nation-states'—was a module that allowed attackers to wipe the victims' disk drive. Despite its inclusion, attackers did not use this capability.

The Odinaff attackers also manipulated the SWIFT messaging system using tactics almost identical to the nation-states'. The malware looked for any strings in the SWIFT messages that included specific details, such as dates and international bank account numbers. When the date and account number in a SWIFT message matched the details associated with a fraudulent transaction, the malware suppressed the message, preventing the bank from discovering the activity or at least delaying it until the funds were already gone.

While no cybersecurity officials have established solid attribution, several clues point to attacker ties to Russia. Strings present in the malware, as well as folder names, were comprised of Cyrillic characters; additionally, some speculated the existence of a relationship between the Odinaff attackers and the Carbanak malware attacks. Carbanak is the tool of choice of a cybercriminal gang, also referred to as Carbanak, that has targeted large corporations for financial gain since at least 2014. The Carbanak gang has been the subject of both media and security reporting due to their high-profile attacks.

The North Korean and Russian-based Odinaff attacks were so similar that, when initially discovered, investigators believed the heist originated from the same North Korean attackers responsible for the previous SWIFT-related attacks. They soon realized that was not the case, but this serves as another example of why investigators cannot let opinion dictate attribution; they must follow the evidence. While the Odinaff attackers were successful—they were one of a few cybercriminal groups to steal money from financial institutions themselves as opposed to their customers—they did not enjoy the same monetary success as nation-state attackers.

Conclusion

Nation-state financial theft wasn't a problem for banks prior to the 21st century. Unfortunately, since 2009, nation-state attackers, including those from Iran, North Korea, and Russia, have conducted attacks that include sabotage, financial theft, or denials of service against banks all over the world. The attacking nations have suffered under sanctions; in turn, these sanctions then motivated the attacks. For example, North Korea and Iran are under sanctions for developing and testing nuclear weapons. The measures in place restrict economic growth in order to pressure both countries to halt their military development of nuclear weapons. Yet the funds obtained through financial theft often supplement this monetary loss, allowing nations to continue building their military power.

In addition to economic motivation, Iran and North Korea conduct attacks to project power in the public eye and to retaliate against alleged U.S.-based or allied cyber operations.⁴⁰ Attacking financial institutions for substantial monetary gains and with large-scale DoS and sabotage attacks sends a message to the government in which the victim banks reside. Other nations, like Russia, have been sanctioned for military activities as well, just not for those involving nuclear weapons. While not discussed in this chapter, Russian attackers usually target financial institutions for retribution purposes and to cause economic turmoil in the targeted nation.

The impact of cyberattacks magnifies when bank customers cannot access their money, resulting in negative media attention for the victim organizations. This media coverage causes embarrassment to banks and often results in a loss of customers who may feel their money is no longer safe. It is plausible that in a country with a weakened economy, this type of attack could impact its overall economic posture.

While these attacks might sound like plots from spy movies, bear in mind that they actually took place, demonstrating the danger that nation-states pose to financial institutions. Nation-state attackers are possibly the most dangerous and impactful threats that financial institutions face today. While nation-state attacks are rare, the monetary loss from a single attack is far greater than that from traditional cyberattacks. For these reasons, organizations need to handle and respond to them differently, as simply blocking or mitigating the initial threat will not stop this type of attacker.

