BLACK & WHITE PAPER

Tech Industry Doppelgangers:

Campaign Innovation in the World of Cybercrime

APRIL 2017

COMMISSIONED BY

About this paper

A Black & White paper is a study based on primary research survey data which assesses the market dynamics of a key enterprise technology segment through the lens of the 'on the ground' experience and opinions of real practitioners – what they are doing, and why they are doing it.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
1411 Broadway
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 207 426 1050

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

Introduction

Decades ago, novelists and Hollywood writers had to exercise their imaginations to find an exciting story in the mundane technology of the time. These days, defenders need a healthy imagination to just to figure out what adversaries might be doing. The events surrounding Stuxnet and Edward Snowden were truly stranger than fiction. People labeled as paranoids and tinfoil-hat wearers prior to these events felt justified afterward. The evolving role of criminal groups and government-sponsored adversaries have changed the game for good.

A unique aspect of software-based attacks is that, when a weapon is used, it is also unavoidably captured by the victim in the process. The result is an arms race where the average criminal could be using attack tools with government-level sophistication mere months after they are first seen in the wild.

As an example, it has been estimated that Stuxnet cost $100 million to develop. Within just a few years, malware samples were discovered that nearly matched Stuxnet's sophistication, but cost as little as $10,000 to develop. At least one of the zero-day vulnerabilities used by Stuxnet is still in common use by malware over six years later. This example of 'reverse hyperinflation' has created the perception that attackers are now advancing at a pace defenders could never catch up to.

Technically, it is true that malicious adversaries benefit from state-sponsored attack tools. However, an increase in the sophistication of malware doesn't equate to an increase in the sophistication of the people using the malware. While the capabilities of malicious adversaries' tools has gone up exponentially, in many ways the people behind criminal campaigns in the digital world aren't much different from average office workers and IT staff. Just as giving an amateur photographer a license to Photoshop doesn't make them a professional, giving an experienced fraudster a copy of Stuxnet doesn't turn them into some sort of cyber-warrior.

While malware is still a serious and unsolved problem, however, it is rarely the entirety of the threat we face against malicious adversaries. For criminals and other types of adversaries to succeed, there is generally a series of events that must take place. **In fact, according to the Verizon's Data Breach Investigations Report, malware is typically involved in less than half of known breaches. Yet security spending tends to be heavily biased toward stopping malware, despite its role as only one piece of the attack campaign.** Perhaps the reason malware is such a focus is because it is both visible and personal – nearly anyone that has used Windows for a significant period of time has suffered from a malware infection. It has become almost normal for some consumers to simply buy a new computer when one gets infected. Economically, it is often cheaper than paying someone to remove the malware. In fact, we've calculated the enterprise cost of malware remediation to be as high as $2,300 per infection.

The truth, however, is that most malware and attack campaigns can be more easily stopped, disrupted or frustrated by better understanding the attacker's tactics, techniques and procedures (TTPs).

What Is an Attack Campaign?

The attack campaign is a detailed plan laid out by the attackers. It is put into effect long before the victim is targeted, regardless of whether the target is your kid's smartphone, a wireless broadband router in Brazil or your company's network. The typical campaign figures out how to monetize an attack based on the tools and resources available to attackers. As much of this campaign as possible is automated, including getting paid. Often, by the time the victim is aware of the attack, the damage is done, the attackers have profited. Even before profits begin to dwindle from a campaign, attackers have moved on, either modifying the existing campaign to evade detection or building an entirely different one.

Figure 1: Attack Campaign Overview

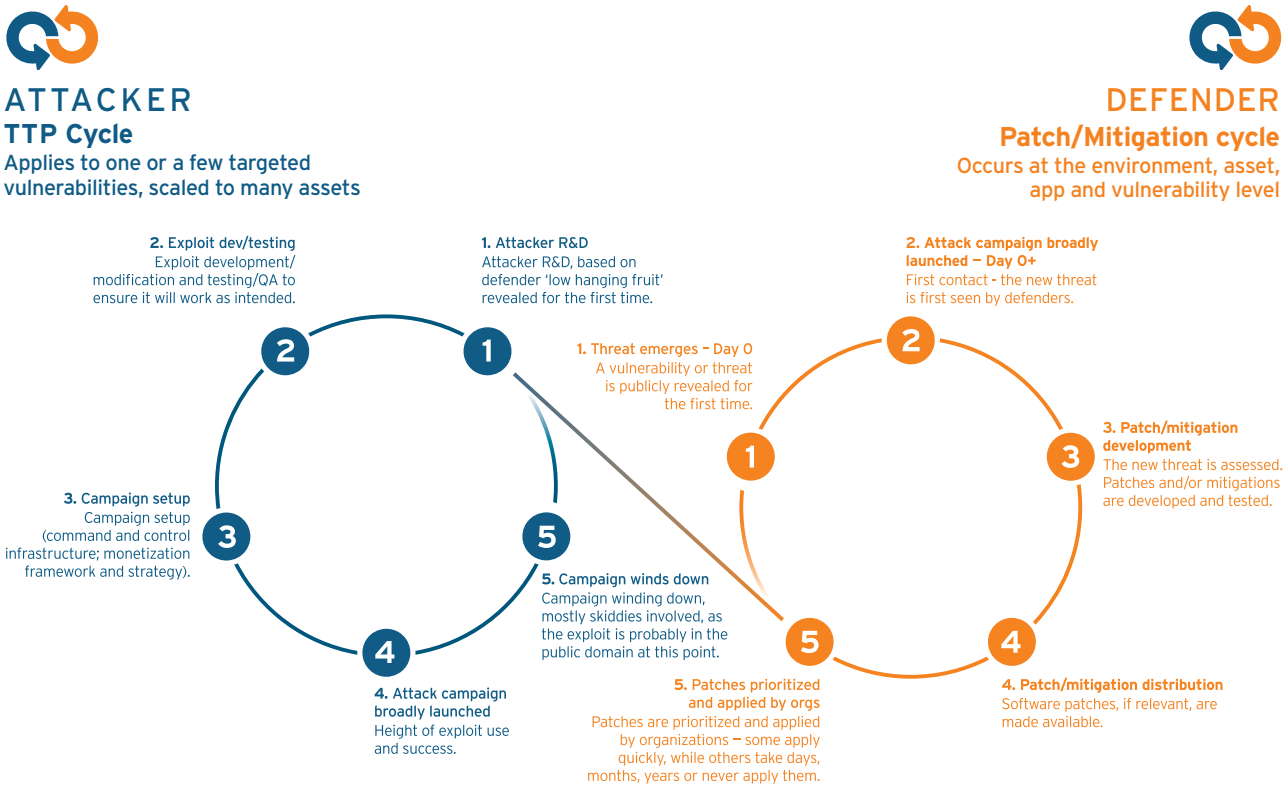

At the atomic level, nearly everything used by the attacker is now disposable, making most threat data and traditional anti-virus techniques almost useless. Industry sources have found that the vast majority of malware (over 95%) is automatically generated to produce unique binaries that are only used once and then discarded. Attack infrastructure, such as domains, IP addresses and servers are also largely disposable to attackers. These are all forensic artifacts located near the bottom of David Bianco's well known 'Pyramid of Pain' This pyramid is used to express both the value of threat intelligence and the difficulty of obtaining it.

Figure 2: Pyramid of Pain

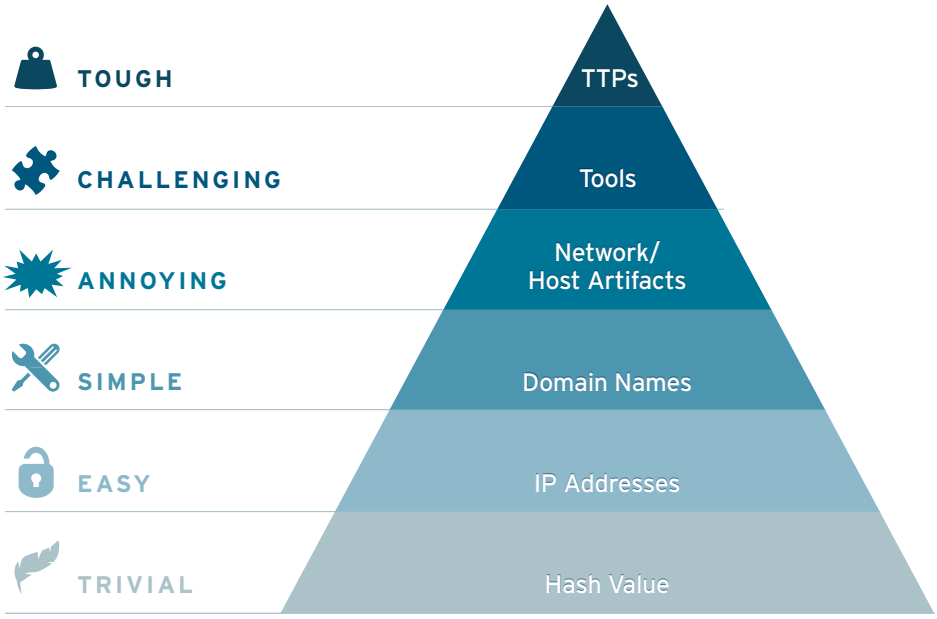

Source: David Bianco

Hooking up a feed of malware hashes to a secure web gateway or antivirus agent is a simple process, but when adversaries started producing unique binaries for each campaign, the approach of looking for known threats lost value almost overnight.

This paper will focus on how defensive tools and strategies can be shifted to address the attackers' TTPs instead of the millions of disposable resources used to compromise businesses and individuals every year. While Bianco's pyramid serves as a useful reference for categorizing and understanding threat data, we disagree with the assertion that threat indicators toward the top of the pyramid are more difficult to obtain or leverage. **There are a comparatively small number of TTPs that tend to resurface again and again in attacks, and many can be addressed through simple configuration changes and other approaches that cost the defender nothing.**

2010 TO PRESENT: THE NEW THREAT LANDSCAPE

The threat landscape has changed again and again over the years, while defenders struggle to cope. The adversaries now include hacktivists and state actors. 'Cyber' is even an official branch of the military for some of the world's larger nations. As nearly everything with a source of power is connected to the internet, and cloud computing goes mainstream, attackers have more targets and opportunities than ever.

The new discipline of developer operations (DevOps), which emphasizes rapid development and deployment, has emerged from changes in software development practices and shifts in thinking on how businesses can and should use technology.

Finally, the targets and dynamics of attacks have changed.

Figure 3: Current Attack Targets

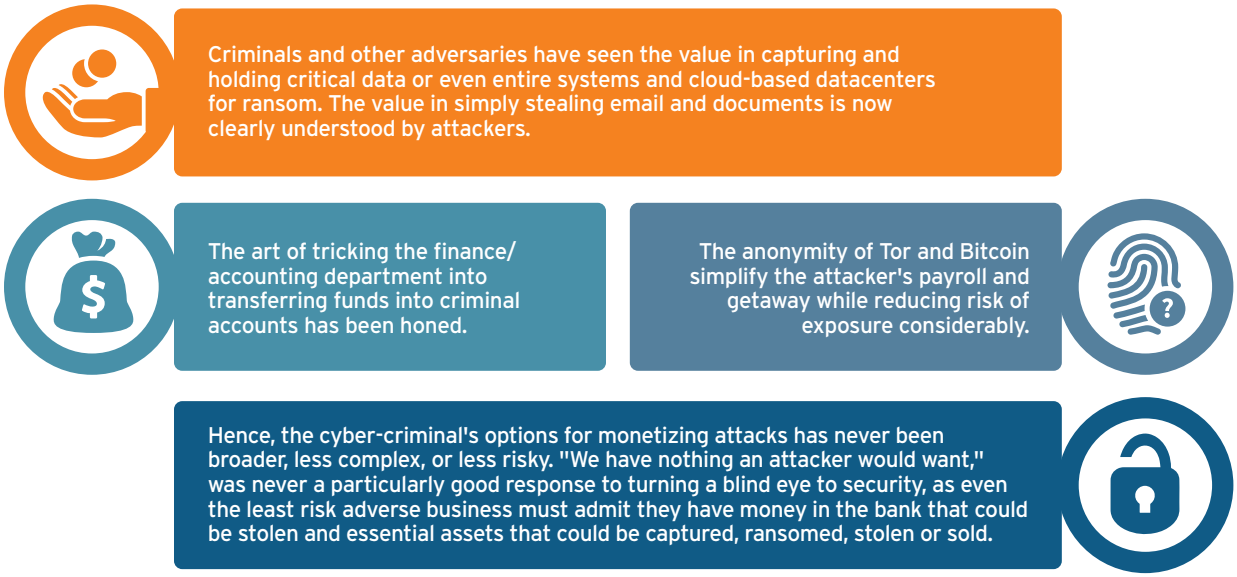

- Criminals and other adversaries have seen the value in capturing and holding critical data or even entire systems and cloud-based datacenters for ransom. The value in simply stealing email and documents is now clearly understood by attackers.
- The art of tricking the finance/accounting department into transferring funds into criminal accounts has been honed.
- The anonymity of Tor and Bitcoin simplify the attacker's payroll and getaway while reducing risk of exposure considerably.
- Hence, the cyber-criminal's options for monetizing attacks has never been broader, less complex, or less risky. "We have nothing an attacker would want," was never a particularly good response to turning a blind eye to security, as even the least risk-averse business must admit they have money in the bank that could be stolen and essential assets that could be captured, ransomed, stolen or sold.

WHY ATTACKERS WIN AND DEFENDERS FAIL: A SUMMARY OF ADVICE AND KNOWLEDGE FROM THE EXPERTS

For this report, we interviewed seven cybersecurity experts with experience in both 'black hat' and 'white hat' methods. Here is what we learned.

KEY ATTACKER STRATEGIES:

- The attacker 'three Rs': Reuse, Recycle and Reinfect. Attackers will use whatever is cheap, free or available, as long as it works.
- Moving beyond financial and identity fraud to strategies that are more profitable and involve less complexity and risk.
- Phishing, ransomware and DDoS are still sure bets for the attacker.
- The KISS principle: Attackers seek to eliminate complexity wherever possible. Attackers follow the path of least resistance!
- Attackers are only as stealthy as they need to be (which isn't very stealthy).

KEY DEFENSIVE STRATEGIES:

- Perform Open Source Intelligence (OSINT) searches, looking for the same errors that attackers are looking for – private keys and cloud credentials in GitHub or other public code repositories. This is low-hanging fruit for attackers.
- Improve visibility in key blind spots: east-west traffic, the endpoint, data and the cloud.
- Defend individuals, not devices.
- Understand that an effective defense usually doesn't deter an attacker, it just forces them down a different path. Predict the attackers' next move before it happens.
- Look for potential attacker TTPs, not just tools or threat indicators
- Identify the 5-10 key software products responsible for the vast majority of infections and disable or mitigate the threat they pose.

STAGES OF AN ATTACK

As mentioned, malware is often only used in part of an attack. To fully understand how different tools and TTPs might be used throughout a campaign, we should understand the stages of a campaign.

Figure 4: Attack Campaign Stages

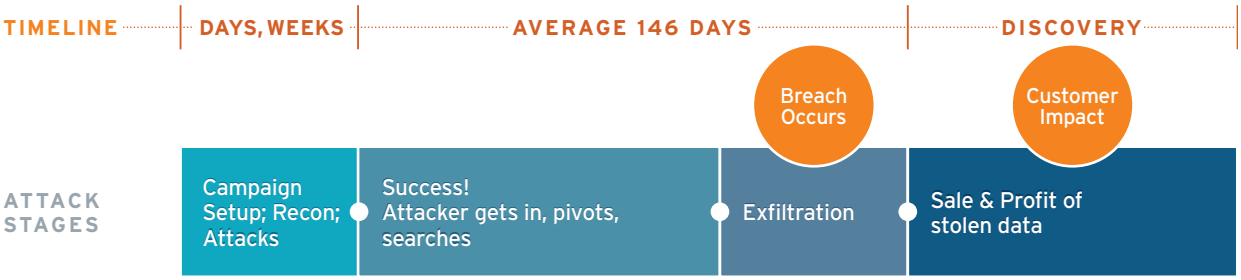

UNDERSTANDING ATTACKER GOALS AND MOTIVATION

High pay, low labor and no patience describe most criminal campaigns. Attackers want the biggest payoff for the least amount of work, and they want it done yesterday. **This profile isn't entirely fair, however – the cyber-criminal world has its fair share of patient, even perfectionist technologists that sometimes meticulously create malware kits, infrastructure and frameworks that see months, if not years, of use across several campaigns. In fact, it is rare that we see attacker campaigns start from scratch.**

More often than not, there are long-compromised systems that are reused to stage payloads, do command and control (C2) duties or just to function as a proxy for phishing or spam emails. Rarely do we see entirely new malware built from scratch. Most malware reuses code and functions, and even infrastructure that's already been built. Sometimes we see the same campaign or a copycat that was previously shut down. **The bad guys have simply moved the infrastructure, tweaked the malware and relaunched from new servers, domains and IPs that aren't hashed or blacklisted.**

As for the goals of attackers, we've seen these criminals move away from stealing payment and identity data. These types of data require a complex and multifaceted business structure to turn stolen data into money in the bad guys' bank accounts. Furthermore, the theft of payment and identity data flooded the market to a point where stolen credit cards became almost worthless, because the supply far exceeded the black-market demand.

We think the combination of complexity and market value are responsible for the current shift back to denial-of-service attacks again and to fully automated extortion (ransomware). In both of these cases, the availability of Bitcoin allows attackers to go from campaign planning to money in the bank in days, rather than weeks or months.

Additionally, the risk of detection and capture are reduced with these types of attacks. **To understand how adversaries will shift in the future, it's essential to think in terms of the lowest effort for the largest payoff; at least in the case of criminal adversaries.**

Reconnaissance

One of the most important aspects of the campaign is how the attacker obtains the initial foothold. The attacker has a few options here – to work from the outside in, or bypass the exterior layers by targeting internal resources. The choice made here likely reflects the attacker's resources, skills and level of patience.

Overwhelmingly, spear-phishing is still the most effective and common tactic used to gain an initial foothold. While defenders tend to see potential targets in terms of lists of corporate-owned assets, attackers are more likely to see a broader landscape, where the target chosen depends on the likelihood of success.

If corporate email systems and devices are heavily defended, an attacker may simply go after a CFO or system administrator's personal email and devices. Attackers know that personal systems are likely to be less diligently defended, but may contain corporate data or the details necessary to access the corporate network.

Attackers will assess all the potential targets and go after the 'lowest-hanging fruit'. A common trend is that we see criminals take advantage of assumptions or poor visibility. The systems that ultimately allow an attacker access are often assumed to be benign or aren't supposed to be in production in the first place. We commonly hear:

- "Oh, that was only supposed to be there for a week of testing, and then removed. That was three years ago."
- "It's just a surveillance camera – what could someone possibly do with that?"
- "That system is in a DMZ, so even if it does get compromised, the rest of the network should be fine."

In some cases, the attacker never needs to touch the company network, since the attacker's goal has already been accidentally leaked by an employee. **In the age of cloud storage with one-click sharing, public code repositories and crowd-sourced technical help sites, what the attacker is after may already be sitting somewhere on the internet, indexed by Google and just waiting for someone to find it and take advantage.**

Criminals will search GitHub for AWS and SSH keys. They will check vendor forums and Stack Exchange for naive administrators posting detailed configuration files in hope of assistance with troubleshooting. Think about it – if credential theft can occur before even touching the victim's network, what chance do they have of detecting the attacker when they look no different than a perfectly valid employee or session?

Figure 5: Defensive Strategies by Campaign Stage

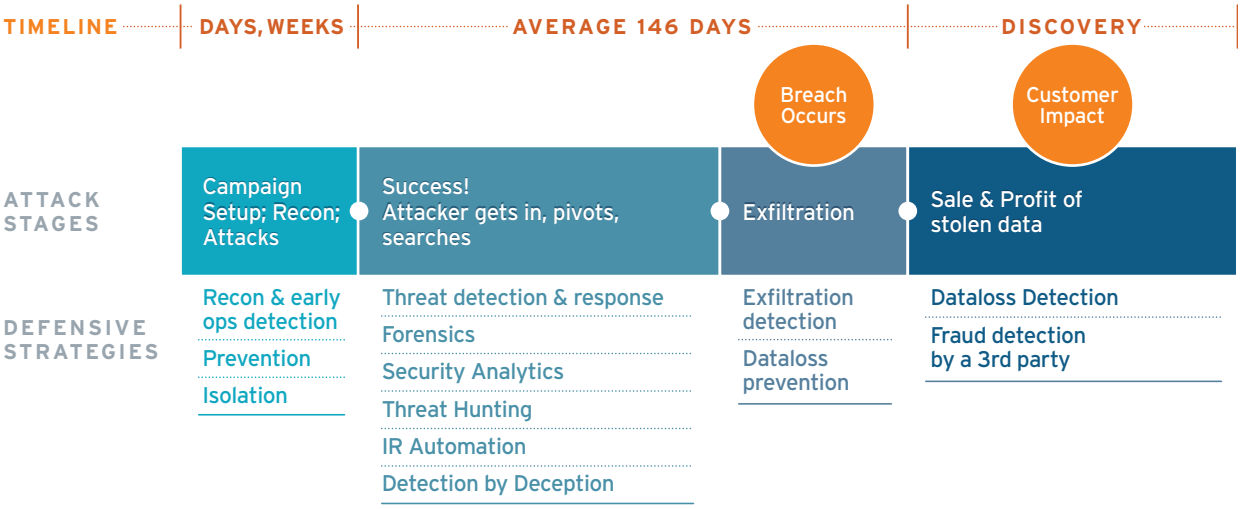

**Dwell time data from 2016 M-Trends report from Mandiant, a subsidiary of FireEye*

DEFENSIVE STRATEGIES:

Defend the individual, not just their corporate accounts and assets.

Don't downplay vulnerable devices just because they don't directly represent an attacker's goal, or 'seem harmless'. Any network-connected system could represent the low-hanging fruit an attacker needs as a stepping stone or launching point for an attack.

Be aware of your network footprint

- Use https://bgp.he.net, Robtex and DNS records to understand the extent of your company's external footprint.
- Use passive services like Shodan or Censys to see how your infrastructure might look to an attacker.
- Use active SSL (https://ssllabs.com) port and vulnerability scans to perform more thorough scans of open services that could represent an entry point for attack.
- Use threat intelligence and blacklist consolidation tools like PassiveTotal and Maltego to look for signs that your company's resources have been used in malicious campaigns (as proxies, sending spam, participating in DDoS attacks, for example).

Look for signs that sensitive corporate data or strategic information has been posted or leaked to the public internet.

- Use search engines like Bing and Google.
- Look on Pastebin sites, Twitter, and forums for signs that stolen data belonging to your company has been leaked.

Ensure every public-facing system is assigned to an owner and is actively maintained — abandoned or forgotten systems are a common source of low-hanging fruit for attackers.

Pay attention to current attacker trends. Attackers often don't get things right the first time, and we are privy to their less successful 'test runs.' Malware targeting IoT devices and Windows ransomware were discovered and studied by researchers years before attackers got the formula right, and these threats began to cause serious damage.

Initial Foothold

This is the step of an attack that's the most feared, so many organizations put the majority of their budget against it. **Attackers have little trouble stepping around these defenses, unfortunately. This is due to these controls and technologies being overwhelmingly focused on prevention.**

That means these systems are looking for threats they expect – threats that are known. That is a huge mistake because:

- Malware authors and the attacker community learned how to get around these types of defenses nearly a decade ago and are well-versed in it today. For example, if a malware sandbox aims to intercept all 32-bit Windows executables, the attackers hide the executable in a Java JAR file. If the sandbox starts inspecting JAR files, the attacker creates office documents with malicious macros and instructions directing employees to enable macros when they are disabled by default.
- If all else fails, attackers use an encrypted session, since most businesses aren't inspecting traffic encrypted with SSL/TLS. Emerging services that offer free SSL certificates make it simple for attackers to encrypt sessions. Traditional network security defenses are infinitely evadable. It is a game of leapfrog that just doesn't end.

The other issue is the assumption that the attacker must attack from the outside in. The reality is that many attacks work from the inside out, through links and attachments in emails, and drive-by website attacks aimed at vulnerabilities in web browsers and plugins.

Attackers will also target users outside the relative safety of the corporate perimeter if they can. It is common for malware to use several stages or components. The initial stage doesn't look like typical malware, because its job is just to obtain the initial foothold and download the actual malware payload that performs the dirty work. This 'downloader' component has an easier time getting past network defenses – the next step is where things get interesting. The use of DNS as a channel for malware communication and even data exfiltration has long been a standard approach.

Attackers have a tendency to share and reuse TTPs. If it keeps working, they'll keep using it. They even have detailed manuals, IT support services, ad tech experts and graphic designers available for building a campaign. Attackers compile, share and hand down instructions, manuals and intelligence.

The IP addresses of honeypots and crawlers belonging to researchers and security vendors are compiled and avoided. Malicious websites will even serve up a benign version of a malicious site to visitors that appear to be investigators rather than victims.

The magic of automation also comes into play on the attacker side. Automated attack platforms can auto-generate malware guaranteed to bypass anti-malware products. The major anti-malware vendors are monitored, and as soon as a particular malware sample is caught, flagged and analyzed, the platform knows to automatically revoke its use and replace it with a fresh sample. Defenders, through information-sharing groups, are increasingly realizing the advantage of sharing intelligence and strategies, but still have a long way to go before they catch up with a well-organized criminal outfit.

DEFENSIVE STRATEGIES:

Focus network defenses on detecting TTPs instead of tools or threat indicators.

Integrate network defenses with endpoint defenses and other security tools. By correlating more contextual data, it's possible to make more accurate decisions about whether the combination of behavior from two perspectives is legitimate and normal or anomalous and malicious.

There's a Bull in Your China Shop: When External Attacker Becomes Inside Threat

Once on the inside, the traditional 'perimeter heavy' business is challenged to detect internal operations of attackers. This isn't because attackers are particularly stealthy, although they can be.

Attackers are generally as stealthy as they need to be, which is all-too-often not stealthy at all. Studies of large, well-known breaches show that not only are attackers often noisy, they actually often get detected by security tools!

Why don't staff see these alerts? The failure occurs due to the all-too-common effect of alert fatigue. **Even though tools detect the attackers' tools or TTPs, security systems spew out so many false positives and less important (but still valid) events that the truly important messages get buried in all the noise.** The result is that average dwell times – the time an attacker remains undetected inside an organization's systems – number in the hundreds of days. Still, as security products are updated to detect and defend against known TTPs, attackers are forced to pivot and change methods.

The most critical point for an attacker to evade is the perimeter. With the majority of InfoSec budgets focused on defending the perimeter between the public Internet and the internal corporate network, this is the primary challenge for invaders to break through. So, how do they? Well, they generally don't break through; they slip through. With Internet-facing network services being better and better defended these days, attacks against the internal network focus on something that's already there – employees.

Emails with malicious attachments and links still get through and still work. Mass emailed threats are easier to recognize and block, but spear-phishing is still a very reliable form of attack. Consider also that most employees check personal mail while at work, giving attackers twice the opportunity, if the personal addresses can be uncovered. This is not a difficult task for the professional criminal. In fact, most criminal groups already have access to vast spam databases that contain email addresses and other personal details. These databases can be mined to match personal addresses with corporate ones.

The web browser is also still a target, although it's one that has been almost obsessively improved to defend against attacks in recent years. Hundreds of thousands of dollars in bounties have been paid out to white-hat hackers that can successfully bypass the defenses of browsers like Google's Chrome. Microsoft chose to build a new browser, Edge, from scratch with security as a key feature, deemphasizing the use of Internet Explorer unless legacy business needs require it.

How is it that many attacks still come through the browser, then? Browser plugins. There's a good reason why Google Chrome comes with a native flash plugin and PDF viewer. Adobe Acrobat and Flash are two of the favorite products for attackers to exploit. In addition to these, Microsoft Silverlight, the Java browser plugin and Microsoft Word are often targeted by malware.

This move away from targeting the web browser itself and toward adjacent software is a systemic issue in how we defend our infrastructure. We tend to think one step ahead, addressing only the current target of attackers and not the next target. We need to stop playing leapfrog and start playing chess. Had we planned several steps ahead, it might have been obvious that once we started focusing on regularly patching and better defending the operating system and browser, attackers would go after all the other software installed on these systems.

Until fairly recently, most businesses and even some tools haven't made it a priority to assess third-party software. The truth, however, is that over 90% of the malware causing headaches targets less than 10 products:

- Microsoft Office (on Windows)
- Oracle Java (typically, the browser plugin)
- Adobe Flash
- Microsoft Silverlight
- Microsoft Internet Explorer
- Adobe Reader/Acrobat

What's more is that most of the exploits currently in use are years old. One of the vulnerabilities used by Stuxnet (CVE-2010-2568) is still in active use by some malware, and patches were released by Microsoft for versions of Windows from Windows XP on up through Windows Server 2008. Despite patches being released over six years ago for all Windows platforms, attackers still seem to be having success with it.

DEFENSIVE STRATEGIES:

More than just the operating system and its components need to be patched and/or protected. **Determine if all third-party software is actively needed and used.** For example, we find in many cases that businesses might need Java to run local applications, but have no need for the Java browser plugin to be loaded. In some cases, we've heard of 100% of malware infections being prevented by simply systematically removing the Java browser plugin after verifying employees did not need it or use it.

Industry reports often have helpful summaries listing the software, vulnerabilities and exploits used by attackers. In particular, the Microsoft Security Intelligence Report provides a pragmatic and actionable listing of the current threats.

Pivoting, Internal Recon and Credential Theft

Attackers have increasingly favored taking advantage of tools already in place. In the early days of hacking, adversaries were often amateurs and attempted to build in everything they thought they would need – the old BO trojan, for example, had broad remote system management capabilities built in. **Now that many attackers are professionals, it is common that they have basic IT skills and perhaps have even worked in or currently work in legitimate IT roles. The result is an adversary more likely to solve a problem like an IT administrator, not like Hollywood's movie depiction of a hacker.**

A trojan like BO is simply redundant in an environment where full remote functionality is built into most products already – this is a minimum requirement for any environment where thousands or tens of thousands of systems must be administered by only a few IT staff.

While malware still often plays a role in part of a campaign, we nearly always see the use of common IT tools in a modern-day campaign. Windows Management Instrumentation (WMI) allows for remote changes to be made to Windows servers, desktops and laptops. An attacker with stolen credentials can easily take advantage of this, often without setting off alarms. The Simple Network Management Protocol (SNMP) is the WMI equivalent for many network devices, and is one of the most common vulnerabilities in businesses, since its default values are often easily guessable. Microsoft's SysInternals suite of tools are incredibly useful for an IT administrator, and for an attacker – often these tools are already present on servers and workstations for the convenience of the helpdesk and admins. Microsoft's Powershell, Active Directory and Group Policy framework also offer an attacker a broad set of automated tools that are conveniently available once valid credentials have been created or stolen.

Sometimes-Malware? Sort-of-Malware?

What about when legitimate software is used for malicious purposes? Clearly, we can't treat every administrative tool like a piece of malware, because the IT staff depend on these tools. Similarly, not all software categorized as malware poses a risk to the business.

In one example, a considerable investment was spent on responding to the presumed threat that the infamous Conficker malware could have posed. In the end, through five variants and 163 days of activity, only the final version of Conficker downloaded scareware onto victims' PCs for a period of only 26 days. Afterward, it removed the scareware and has remained dormant ever since. Conficker is likely still infecting machines today, though it has no malicious payload or instructions to carry out.

Regarding a more recent issue, opinions vary on whether government organizations and law enforcement agencies should be allowed to use what is arguably malware against citizens suspected of crimes. The answer is that we shouldn't be looking at software in isolation. Instead, we should be looking at how software behaves. A hammer isn't inherently good or bad, but one person swinging a hammer at another person is a malicious behavior associated with that tool.

DEFENSIVE STRATEGIES:

Instead of just looking for malware, we must look for anomalies on systems. These anomalies might be the result of a malware infection, or an attacker that had no need of malware in the first place.

Whatever the scenario, systems should be configured to detect malicious or suspicious use of legitimate software.

Also, create a separate category for legitimate software that is not used by the IT team. For example, if the standard for remote PC access is Microsoft's built-in remote desktop tools, the use of the third-party products such as Dameware or TeamViewer should be suspicious, even though these products might be normal in other environments. The difficulty here is that this kind of anomaly detection has to be custom configured for each company's environment, since each company has a unique baseline for what 'normal' looks like.

Communication With the Outside

A common preference we've seen among attackers is the use of free software and products wherever possible. In one novel example, malware used cloud storage utilities (aka 'file sync and share' [FSS] software) for command and control (C2) and exfiltration. These Cloud FSS services used WebDav to synchronize, saving the attacker the trouble of having to create that function. Additionally, leveraging a legitimate piece of software helps the attacker hide C2 communication from easy detection. In both cases that used FSS software, communications and data were encrypted before transmitting in both directions, keeping network or perimeter security devices from peering into the data coming in or going out.

While some attackers change tactics, others don't, but still succeed. Why? Often because they succeed in hiding among valid business activities. Other times, because many best practices still have yet to become common practices. We've seen recent major breaches succeed in sending off gigabytes of unencrypted data over regular old FTP to servers in countries where the victim had no legitimate customers. Either one or many layers in the modern secure enterprise should have caught or at least seen this activity:

- GeoIP correlation at the perimeter should have been able to detect or block transfers to a country that the victim never communicates with.
- Detection of a large outbound transfer to a source that has never before received large data transfers should have raised a red flag.
- Use of FTP, an old and outdated protocol, should alone raise a red flag.
- DLP should have been able to detect or block sensitive customer data transferred by FTP.

DEFENSIVE STRATEGIES:

Clinically, the fear of getting rid of things is termed disposophobia. It isn't OCD or phobias that are responsible for the commonality of alert fatigue in NOCs and SOCs, however. InfoSec teams particularly have good reason for not wanting to get rid of alerts, logs or other stored events. Often, it isn't clear which messages could be critically important until an investigation is already under way. Even the most mundane event – an 'up event' on a switch port, noting an Ethernet link – might be important in an insider threat case.

This 'noise,' however, often cripples and blinds a business to the true threats, but the fear of missing something important makes many security teams hesitant to filter alerts more aggressively.

The key is not to get rid of events, but to prioritize by separating red flags from yellow flags. Red flags are alerts that have little to no chance of false positives, while yellow flags are the 'maybes' that make up all the noise produced by most security systems.

Further prioritize and tune security alerts by simulating attacker TTPs and then noting the artifacts that allow detection, and deemphasizing the artifacts and alerts that don't contribute to detection.

One promising approach is anomaly detection. Without systems to learn what normal traffic patterns look like for a business, these sorts of anomalies would be invisible. The challenge here is that most security products are developed with a 'one size fits all' approach to design. Detecting attacks by looking for anomalies requires an understanding of what 'normal' looks like, and a product or staff need time to build that profile.

EXAMPLE:

You've baselined your company's network traffic, resulting in a clear profile of what's 'normal.' A sudden change in FTP traffic from 0 megabytes per day to 5 gigabytes per day from a server with no requirement to talk to the outside world is likely to be malicious activity.

Breaking the (Kill)Chain

As mentioned, the success of attackers in the last 4-5 years has resulted in a black market with more payment data and identities than necessary to meet current demand. The effects of Edward Snowden, Julian Assange and state-sponsored campaigns have had an unexpected effect on attacker strategy and behavior. After the events of the last five years, we've seen attackers realize the value of information beyond what can easily be sold in a forum, and we've seen the adoption of more sophisticated tools and techniques.

Going through the full kill chain requires a high amount of skill, patience and ingenuity. While we've always seen interest in campaigns that don't follow the traditional pattern into a corporate network, the incentives to seek safer, quicker and easier paths to a paycheck are now clear.

- SQL Injection (SQLi) is a devastating attack that's still often used in corporate breaches. It isn't new by any stretch, but we keep seeing it used and successful. It's the 'checking the sun visor for car keys' of cybercrime – it doesn't work as often anymore, but it is such an easy paycheck, attackers would be crazy not to check for it. Think about it – only one vulnerability to exploit and no need for malware, pivoting or exfiltration strategies. Once a SQL injection flaw is discovered, the entire attack campaign occurs in one step.
- Ransomware fully automates the attack in a way that puts money directly into accounts – no need to launder it through a mule or through selling merchandise bought with stolen credit cards. It cuts down on time-to-payment for the attacker, as well as the risk of getting caught.
- Attacking data in the cloud: A company may have a formidable firewall, but if the attacker doesn't have to cross it, they won't. Nearly all businesses have some data in the cloud now, and it has been commonplace for years for attackers to go after it. We've seen criminals go after Salesforce, downloading a customer's entire dataset, if possible.
- Stealing credentials without touching the user or the network: It has become standard procedure for attackers to obtain large password leaks, or perform the compromise themselves to obtain large lists of email addresses, usernames and passwords. The email addresses can be added to spam lists, and thanks to rampant password reuse, the credentials can potentially be used anywhere. We've especially seen this activity after the Adobe and LinkedIn breaches. Since the email address has almost universally become the username in the username/password credential pairing, it becomes trivial to test them against other services. Scripts can automatically test these credentials on popular social media and banking sites (going after consumers), or can test them against corporate resources and enterprise-focused SaaS apps like Office 365. We know of several major breaches that allegedly occurred due to password reuse extending to corporate resources.
- Going after new types of data: While we didn't see any attempt to monetize the emails leaked in the Sony Pictures and Ashley Madison breaches, company correspondence is often one of the most valuable types of data in any organization. In addition, it tends to be centralized (on the company email server, or in the cloud) and highly distributed (local copies of email on employee systems and in PST files, which could be in backups or even on file servers), making it very difficult to effectively protect. We are always looking for new ways that ransomware might be utilized beyond encrypting files, and threatening to reveal sensitive email content is high on that list.
- DDoS attacks: along with SQLi, DDoS is one that never goes away, because it just keeps working. The latest twist in the long trend of DDoS attacks are the recent IoT-based botnets assembled by simply scanning the Internet for open TELNET services and attempting to log in with a short list of common default credentials.
  - » This has resulted in botnets like Mirai, which have demonstrated attack capabilities over 1Tbps for the first time, thrusting us into an age where DDoS is no longer just a tool for taking down a website. With this much raw power and by strategically selecting multiple targets, we're looking at a tool that can take down entire organizations, cities, or even small countries with too few paths to the Internet to effectively route around a large-enough attack.

» There is evidence that Mirai is open for business, to be rented for a price. With significant investment in the malware (over 34 Linux ELF binary variants in the past four years), the framework and the infrastructure, we doubt criminals will stop targeting vulnerable Linux systems attached to the internet anytime soon.

The bottom line is that going through the full kill chain is difficult, and professional criminals aren't going to put any unnecessary roadblocks between themselves and payday. If simple attacks work, we'll see more simple attacks. If we make things more difficult, we'll see more complex attacks rise again in popularity.

DEFENSIVE STRATEGIES:

Obtain lists of leaked credentials where possible (obtaining passwords isn't necessary), and compare leaked usernames against employee usernames and email addresses. Require anyone with matches to reset passwords as soon as possible.

Multifactor authentication (MFA or 2FA) is an easy win where it doesn't over complicate user workflows. In cases where credentials are leaked, MFA should be effective in preventing the leaked credentials from being successfully abused.

To prevent SQL injection, sanitize input and prevent free-form SQL input from application service accounts by using stored procedures and triggers. The Open Web Application Security Project (OWASP) has much more detailed recommendations and instructions on how to find and remediate SQLi flaws.

METHODOLOGY

The goal for this paper was a challenging one: "Providing real-world insights into adversary attack campaign strategies, planning and execution," was the guidance given.

We interviewed several white-hats experienced in dealing with the black-hat side of things from various roles: incident responders, malware researchers, penetration testers and all-around cybersecurity veterans. In addition, we've leveraged our own experiences and what we've learned from various defenders, vendors, researchers and even criminals over the years.

The original plan was to gather stories of criminal campaigns and present them up as short vignettes. We quickly found that all the truly good stories had either already been told, or they hadn't been told because our interviewees were sworn to silence where these stories were concerned. Besides, taking the storytelling route would limit the number of examples we could provide and the specific TTPs we could share.

These individuals collectively have over 100 years of security experience, have investigated hundreds of incidents and given hundreds of talks on the subject.