

Phishing – Baiting the Hook

Table of Contents

- 1 Letter from the Editor
- 2 Overview
- 3 Guest Essay: The Data Science Gambit
by Alex Pinto, Verizon
- 5 Akamai Research
 - 6 Phishing in a Nutshell
 - 9 Phishing Kit Breakdown
 - 12 Special: Using Akamai to Defend Akamai
 - 14 Phishing Software Development Lifecycle
 - 17 Phishing as a Service
- 20 Looking Forward
- 22 Appendix: Methodologies
- 24 Credits

20

22

24

Looking Forward

Appendix: Methodologies

Credits

24

Credits

24

Credits

24

Credits

24

Credits

Letter from the Editor

Martin McKeay
Editorial Director

Data science is hard. But data itself is malleable and open to interpretation.

Before I started learning about statistics and data science, I thought data was data and that was all there was to it. It was a naive view based on my understanding of the end-product, rather than the processes and work that went into creating a report. I've since learned that a finished report is only the tip of the iceberg.

Almost any data scientist will tell you that a significant part of their time is spent on cleaning data to prepare it for analysis. Whether we're talking about log files, event histories, or any other type of data, issues always need to be resolved before the work can begin. In some cases, it's missing data. In others it's outliers: events that skew the data or were misreported. Or maybe multiple data sets need to be brought together to form the data you want to analyze. No matter the cause, it takes time and work to ensure the most accurate data is available.

The data we work with to create the State of the Internet / Security report poses all those issues and more. We incorporate disparate data from across solutions that represent some of the same information in very different ways. How the data was recorded, where it was logged, and how many levels of abstraction have happened between the original capture and the final use all influence the data's cleanliness.

We've been reporting on Distributed Denial of Service (DDoS) attempts and application attacks since the report was first published. Over the past three years, we've reached deeper into our organization to gather data on credential stuffing, phishing, and bots, just to name a few topics.

Looking forward, we expect we'll continue to find better data to highlight cyberattacks on a scale few other organizations experience. Some of the data we've been using will be seen less often as we move to new collection methods and cleaner data. We may lose some long-term indicators in the process, and that is part of the trade-off that needs to be considered in any shift.

Data science is hard, but it's worth the effort. Each data set can be manipulated to make the most out of its strengths and compensate for its weaknesses. After that, we can prepare the best analysis possible to tell you a story about the attacks we see.

Overview

Phishing is a long-term, socially based problem, impacting multiple market segments and people from all walks of life each day. This edition of the State of the Internet / Security report centers on phishing and its impact to the retail sector.

In retail, phishing victims are consumers. In addition to other methods, criminals use phishing to target the retail industry by masquerading as popular brands and retail outlets. The individuals who fall for phishing scams by submitting information, or those who inadvertently install malicious applications, are the same people who contribute to a billion-dollar retail economy worldwide. They're a key part of the phishing lifecycle.

Phishing is often just one part of a larger attack against an organization. One of the most chilling examples of this in the retail sector is the late 2013 attack against Target Corporation, which was [first reported by journalist Brian Krebs](#).

Considered to be one of the largest retail attacks in history, Target and millions of consumers were victimized by criminals, allegedly based in Ukraine, who compromised payment card details

for 40 million accounts between November 27 and December 15, 2013. The attack started with a phishing campaign against an HVAC vendor used by Target, and ultimately led to malware being installed on the vendor's systems, which compromised VPN credentials.

From the first line of code used to develop a phishing kit, to the full-scale suites of phishing services sold by criminals, the phishing economy continues to exist, despite the growing wave of awareness training programs and endpoint defenses. It isn't that defenders are losing the fight, it's that the criminals refuse to give up on one of their core markets – which fuels additional scams and attacks such as identity theft, retail fraud, and credential abuse.

In order to adapt to increasing defenses, phishing has evolved from being an email-based attack to one that now includes mobile devices and social media. This evolution is leveraging the world's increasingly connected existence as a means of rapid propagation. This means that criminals now have more options when it comes to targeting their victims.

Phishing at a Glance / TL;DR

- More than **60%** of the phishing kits monitored by Akamai were active for only 20 days or less
- **High tech is the top industry** targeted by phishing, according to Akamai's data, followed by finance, online retail, and media
- According to Akamai's monitoring, Microsoft, PayPal, DHL, Dropbox, DocuSign, and LinkedIn are all top targets when it comes to phishing

The Data Science Gambit

Alex Pinto

Verizon Business Group

Team Leader

Information Security has a fever, and the only prescription is data!

Over the years, I've heard many practitioners say that Information Security is a profession in its infancy. I believe that this anxiety about our profession boils down to a fear of the unknown and a lack of data. We're playing chess in the dark, unable to plan our next move, let alone see our endgame.

Our profession's track record for publicly sharing data about the incidents and threats that individual organizations experience is weak, historically speaking. This leads to a lack of understanding of our controls, their effectiveness, and which technologies are actually doing the most to protect us. While this isn't the place to discuss vendor incentives, Ian Grigg's "The Market for Silver Bullets"¹ is an excellent primer on the topic. It's as relevant today as it was in 2008.

Secrecy, driven by a desire to maintain security through obscurity and a desire to keep breaches private, has played a significant role in keeping us in the dark. The introduction of California's SB 1386 in 2002² was the first U.S. law that requires businesses to disclose breaches. Today, almost every U.S. state has adopted and extended similar breach notification laws. The European Union is taking breach notification even further and driving significant changes in the industry with the General Data Protection Regulation.

The breach notifications these laws require are a boon to our understanding of security, although they require organizing and normalizing data for analysis that may be beyond the scope and/or skill of the average practitioner. Researchers at Verizon recognized the importance of using this type of data and making the analysis available to the community, and in 2008 the Data Breach Investigation Report (DBIR) was created. The DBIR team also saw that breach data was cluttered and messy, which led them to create the VERIS framework,³ a common standard to create a more organized and coherent view of the data.

“ *The game of chess is like a sword fight. You must think first before you move.*”

– *Shaolin & Wu-Tang*

While this didn't immediately create a bright future flush with data transparency, it inspired other organizations to create data-driven security reports. The Akamai State of the Internet / Security Report was created in 2014 to be another voice in the data-starved deserts of information sharing. Prior to my involvement in the DBIR, I saw both teams lead and inspire other companies to share data and promote analysis as a benefit to the entire community. Now that I'm directly involved, I see the work that goes on behind the scenes, and it's amazing the amount of effort that goes into making these reports happen.

Bigger and bolder data sharing is vital for our industry to grow, learn, and better protect the organizations who entrust our profession to do so. We need to do more to share with each other – whether it takes the form of business-to-business data in a secure and privacy-preserving way, or public-facing reporting by organizations like Verizon and Akamai. We need a better understanding of the threats we face as individual organizations, and the research and analysis public reporting provides is one of the best resources we have.

In my opinion, data science and public reports are a vital part of how we move forward in this industry. As more organizations share their data, their intelligence, and their analysis, the better our profession will be as a whole. Unfortunately, too many organizations still hide the underlying data and the purpose of their reporting. As an industry, we should embrace the organizations willing to be forthcoming about their data, methodology, and analysis concerning security reports.

As you read this report, and hopefully [Verizon's DBIR](#), think about how we can be better data-sharing neighbors. Reach out to your industry Information Sharing and Analysis Center (ISAC) and other reputable reporting organizations. Help us shine a light on the darkness to reveal the chessboard, so we can make our next move.

Alex Pinto is a Distinguished Engineer of the Security Solutions Group at Verizon Enterprise Services, currently managing the Verizon Security Research team, which is responsible for the Verizon Data Breach Investigations Report (DBIR).

Alex has over 20 years of experience in building security solutions and products and the last 6 of those years have been solely dedicated to the application of data science techniques on cybersecurity.

Akamai Research

Phishing in a Nutshell

Phishing is a social-based attack. Sometimes phishing may be referenced as a sophisticated attack, but it isn't. These attacks focus on human nature more than a software vulnerability or system exploit.

Phishing is often lumped into the context of spam. But since phishing is a dangerous type of spam with serious consequences, it isn't something to ignore or dismiss as a harmless attack type.

Criminals running phishing attacks will prey upon natural human emotions, as well as the inherent trust people place in total strangers. When it comes to basic communication, either via phone call, email, text message, or social media post, no one expects to be scammed. Why would they? But this natural presumption that things are OK is one of the key elements of a phishing attack.

As phishing evolved over the years, process injection was added to the baseline attack. Process injection works just the way it sounds. The phishing attack targets the workflow, or process, used by an individual victim, enabling a higher degree of success, as well as the possibility for the scam to remain undetected longer. You see this method deployed during business email compromise (BEC) attacks, including those centered on wire transfers and tax-related scams.

Phishing attacks require two things: a lure and a landing. There are times when the lure is also the landing, so the technical elements of an attack depend on the phishing campaign itself, the scope, and the targets.

A lure gets the victim's attention, by way of a warning, an urgent request, or some other message invoking a sense of alarm or concern. Once the lure works, the victim needs to land, and that is where the final phase of the attack happens.

The landing can be anything, including malicious attachments or links, a perfect clone of a bank's website, a retail portal, or a simple form [requesting information in exchange for some type of prize or reward](#).

Outside of malicious payloads, the landing phase could target just usernames and passwords, but sometimes this phase also targets personally identifiable information (PII) and other sensitive corporate or financial information. No matter what information is collected, you can be sure it has value to the criminal.

According to the 2019 Verizon DBIR, 32% of all breaches involved phishing, and such attacks were present in 78% of cyber-espionage incidents.

Most landing elements in any given phishing attack involve a platform, better known as a kit. For this report we'll focus on phishing kits. As a whole, phishing kits are anything but consistent. They change depending on the target, purpose, and criminal's intent. Some phishing kits are highly advanced, with custom security features and targeting options, while others consist of nothing but a single HTML form.

Generic Phishing vs. Spear Phishing

Phishing is commonly observed in two types of attacks: generic phishing and spear phishing.

Generic phishing attacks are a numbers game. The criminal blasts their lure out to thousands – sometimes tens of thousands – of potential victims.

The more lures that are delivered, the higher the odds of success. Imagine casting a wide net into a fully stocked pond. The wider the net, the more fish you stand to catch. The same principle applies to generic phishing. When tracking phishing campaigns, it's the generic phishing attacks that make the most noise, so they're the ones often observed by the public.

Botnets are also part of generic phishing attacks, as they're used to deliver messages to a wider pool of potential victims, often avoiding email-based limits imposed by hosts and re-mailer services. Those of you familiar with security history might remember the Storm Worm (or Storm botnet) from 2007. At one point, this botnet [accounted for 8% of all malware](#) on Windows systems.

Storm circulated via email and used current events (e.g., a massive storm in Europe or daily headlines) to lure people to open malicious attachments or click on malicious links. It worked well, and with each new victim the botnet continued to grow. Many classified Storm as a basic spam campaign, but it was phishing in its purest form. The actors behind Storm, who remain unknown to this day, generated emails designed to pique the curiosity of the recipient (lure) and deliver a malicious payload (landing). The simple two-step process was highly effective.

Spear phishing is a targeted phishing attack. Spear phishing attacks usually only target one person or a group (such as a retailer's customer base or a group of activists). Sometimes spear phishing attacks target a whole company and are mistaken for generic phishing attacks at first. What sets them apart from generic attempts, however, are the granular details.

Targeted phishing attacks will leverage open source intelligence (OSINT), or information about the target that exists in the public domain, as well as other layers of not-so-public information to develop lures that are relevant to the victim. These lures can be workflow-related, such as an internal project or group, or crafted to be viewed as "part of the job" to the target. Lures in spear phishing attacks can also be personal, such as selective sales or offers as part of a shopper rewards program.

Spear phishing is commonly seen in nation-state attacks, corporate espionage campaigns, and fraudulent financial attacks in which the ultimate goal isn't basic information gathering, but something more destructive or consequential.

For example, spear phishing can be used to obtain VPN credentials at an organization with the goal



Generic Phishing

A wide net cast over thousands of potential victims



Spear Phishing

A targeted attack, usually against one person or a group

of using those credentials later to attack the larger target. Spear phishing can also be used to deliver backdoors to a dissident's or sales manager's computer, in order to obtain access to documents and contacts. In other cases, spear phishing can be the method deployed to deliver ransomware to a network, where the initial infection starts by convincing someone to open a link or email attachment.

In 2010, an attack known as Operation Aurora targeted at least 30 different companies worldwide. [The Aurora attacks started with spear-phishing emails](#) that were designed to obtain credentials by delivering malicious payloads via zero-day vulnerabilities in Microsoft products. The attackers targeted these organizations to gain visibility into the activities of human rights activists and conduct corporate espionage.

[Akamai was one of the organizations targeted by Aurora.](#) We were fortunate, however, because although a domain administrator account and systems were compromised during the incident, the attackers were seeking specific data that didn't exist.

BEC attacks, in which the goals are purely financial, are also examples of spear phishing attacks. The impact these attacks have on their victims can be devastating. Impersonating an executive or financial manager, the criminal sends a lure to the victim. The lure evokes a sense of urgency and importance, and often requests wire transfers or payroll information.

Because of the people involved, the messages that the criminals send are often seen as routine, or part of the normal workflow. In numerous cases, it isn't uncommon for an executive to make last-minute requests (even urgent ones), so they're processed and addressed as requested. Usually the victim doesn't realize their mistake until later. Similar schemes also target vulnerable populations, in

which, for example, retirees are scammed out of their savings by a criminal posing as a financial planner, bank representative, or government official.

The FBI said that between October 2013 and May 2018, BEC attacks resulted in [worldwide losses of more than \\$12 billion](#), and that figure continues to grow. On August 22, 2019, the U.S. Department of Justice (DOJ) [unsealed a 252-count indictment against 80 people](#) who were all connected to a massive ongoing series of BEC campaigns.

Two of the individuals named in the indictment were involved in schemes "resulting in the fraudulent transfer of at least \$6 million in fraudulently obtained funds – and the overall conspiracy was responsible for the attempted theft of at least an additional \$40 million," [the DOJ said](#).

Wild-Caught Phish

Phishing kits in the wild, those used for both generic phishing and spear phishing, can exist in a number of places.

Sometimes phishing kits are uploaded to a compromised website. When this happens, the attacker has exploited a vulnerability in the website's CMS or on the server itself. Hijacking a domain like this to host a phishing kit takes advantage of the URL's positive reputation and age, which enables the attacker to remain hidden longer. In other instances, the criminal will choose to purchase a domain and hosting package of their own.

Age is important when phishing URLs are considered. Newly created domains – those that are less than a month old – are often flagged as suspicious by security products. Researchers track domain registrations and report domains frequently if they raise any red flags. However,

taking advantage of top-level domain (TLD) sales at a given registrar, criminals buy in bulk and rotate through their collection during a given phishing run. This allows them to keep operating even if one of their domains – or several of them – are taken down or flagged.

In instances like these, a domain that lasts for a few days could yield hundreds of victims, but even those that only last a few hours still return net positive results to the criminal. This is because after the initial outlay of expenses (domains, phishing kits, and perhaps hosting), a criminal only needs a few victims to get their money back. Everything after that is pure profit.

Phishing Kit Breakdown

As mentioned previously, phishing kits are anything but consistent when it comes to their development. However, when you examine phishing kit distribution, including common kits being sold on the darknet, as well as the kits seen by Akamai's zero-day phishing detection engine, a loose pattern emerges.

Kits typically focus on retail and consumer products, banking or finance, and finally gaming. The reason why kits focus on these market segments isn't complex; they're easy to develop and can be used against a wide pool of potential victims. A breakdown of this is provided in Figure 1.

In the high-tech sector, where the bulk of the detected phishing took place, a number of high-profile technology organizations, including those in the retail space, had several kit variants targeting them.

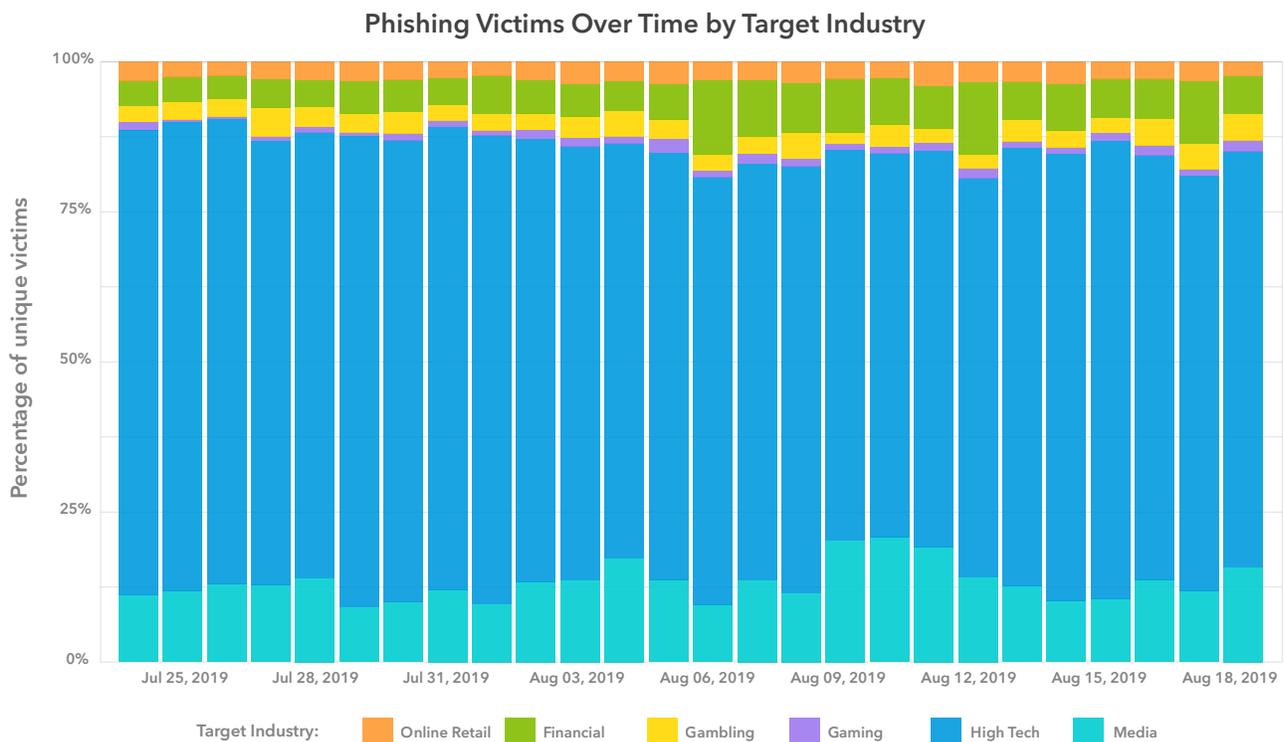


Fig. 1 - Phishing kits may target various brands, but they tend to focus on a few key industries

Although kits may follow a development pattern, the individual kits themselves have multiple variations. The reason for the kit variations comes from a number of factors, including development style, technical enhancements, and evasion methods.

For example, using an observation window of 262 days, Akamai's zero-day phishing detection engine detected 62 different kit variants targeting Microsoft users, which were spread across 3,897 domains. LinkedIn (6 kit variants) and DocuSign (4 kit variants) were also observed across more than 300 domains each.

Kits used in spear phishing attacks could target anything, including the common industries, but they're usually one-off developments that are customized for the task at hand.

Akamai was able to track the lifecycle of each kit from the first time it was observed until the kit stopped triggering our detection rules. Figure 2 shows more than 60% of the kits monitored were active for just

Top Targeted Brands (262 days)

- **Microsoft** - 62 kit variants, 3,897 domains
- **PayPal** - 14 kit variants, 1,669 domains
- **DHL** - 7 kit variants, 1,565 domains
- **Dropbox** - 11 kit variants, 461 domains

20 days or less, which is common among generic phishing attacks. This shortened lifespan is also why criminals constantly develop new evasion techniques that they hope will help keep the kit below the radar.

Cumulative Percentage of Phishing Kits by Days Before Deactivation

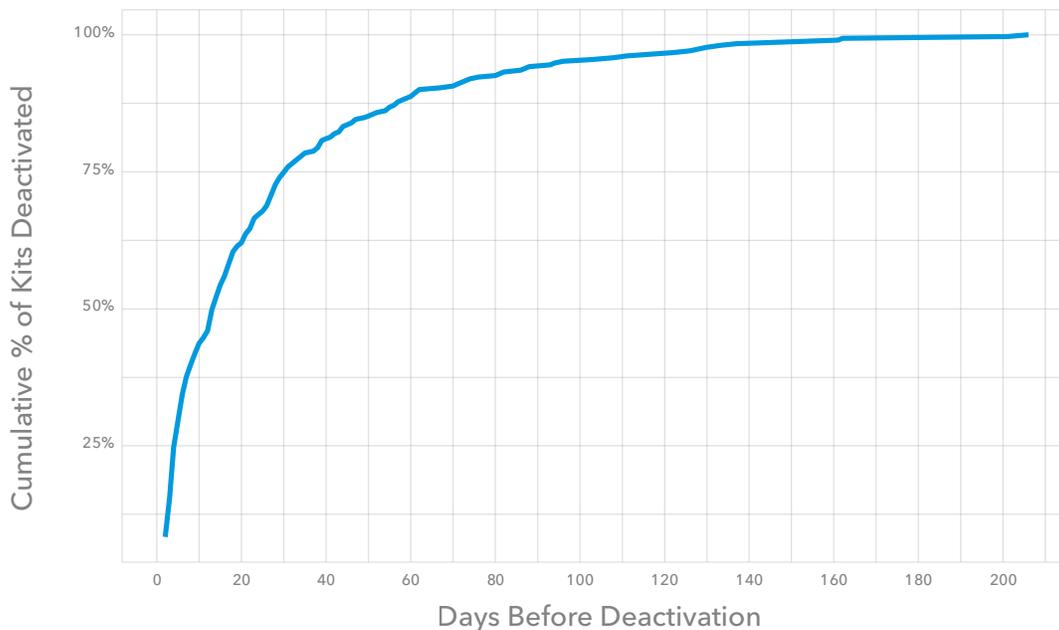


Fig. 2 - More than 60% of all the phishing kits tracked during the 262-day window were active for less than 20 days

Domain Stats by Top Level Domain

TLD	PCT ACTIVE 3 DAYS OR LESS	UNIQUE DOMAINS
.loan	97.95%	1,196,132
.com	96.64%	1,831,417,850
.tk	95.10%	10,637,204
.gq	92.12%	923,345
.ooo	91.66%	147,030
.cf	91.22%	1,263,011
.science	90.85%	69,450
.bid	90.30%	218,995
.ml	89.35%	994,033
.ga	88.83%	1,247,897

Fig. 3 - Looking at the top 10 domains with short lifespans, many of them are commonly associated with phishing, while the .com domains also include those used for botnets

Criminals are in a race against the security teams looking to shut down their operations. Although security teams report phishing URLs regularly, some criminals choose web hosts and domains where those reports are simply ignored. Yet, as the data shows, most kits have a short life, and the window of opportunity for most phishing kits is growing smaller.

In fact, over a 60-day period, Akamai observed more than 2,064,053,300 unique domains commonly associated with malicious activity. Of those, 89% had a lifespan of less than 24 hours, and 94% had a lifespan of less than three days. A breakdown of the TLDs can be seen in Figure 3.

Considering the phishing domains, notable short-lived TLDs such as .gq, .loan, and .tk have a median lifespan of 24 hours and mean lifespan of less than two days. Looking at the data, the availability of cheap name registration on TLDs such as these is a boon to criminals; it makes detection more difficult because the names live in traffic so briefly.

The high number of .com domains with short lifespans can be attributed to names used for botnet traffic, with large numbers of new names used daily (most of which are not registered and so do not resolve).

Using Akamai to Defend Akamai

Being a security company doesn't make Akamai magically immune to phishing. The opposite is probably closer to the truth: Being a security company makes our organization a more appealing target than many enterprises. In February 2019, our CSO, Andy Ellis, reflected on being a target of Operation Aurora,⁴ and tech companies are every bit as much a target of hostile powers now as they were in 2010.

So how does Akamai protect itself against the threat of phishing, one of the most prevalent attacks any organization faces? We use layers of defenses from multiple vendors, including our own product, Enterprise Threat Protector (ETP). Since much of the data in this report is drawn from our global logs and research, we thought it appropriate to share information about what we tell customers about our own experience. These systems represent a subset of our controls, with numerous additional systems in place to protect our enterprise.

While we're not at liberty to name our first layer of phishing defenses publicly, it's a solution we've been using for several years that has historically blocked 94% of the incoming phishing attempts each month. But even the 6% of phishing attempts that reach inboxes is a significant number and a significant threat. The vast majority of attempts, 93%, contained URLs linking to malicious sites. In our experience, these are good results for the first layer of our

defenses. ETP forms the next layer of our defense, using our research and data, augmented with third-party data, to identify malicious domains and block them at the HTTP and DNS level.

Because of the nature of our business, Akamai performs more DNS requests than many businesses of similar size – over 7.4 billion requests every 30 days. ETP has caught 1.2 million malicious requests in the mix, stopping the system or tool from reaching its intended target.

It's important to keep in mind that we identified 2,395 unique phishing threats over the past 30 days and have historically seen more than 4,000 during the holiday months. During the same period, we were able to identify 120 different campaigns, meaning the common elements in the phishing attempts gave us significant confidence they were created by the same tools, even if they had some variation in the content.

As seen in Figure 4, engineering teams at Akamai were the target of nearly 27% of phishing attempts. Given that some engineering team members have responsibilities that expose their email addresses to the public, it isn't surprising to see them harvested and used. Similarly, a large amount of phishing attempts (22%) targeted public accounts, such as "help@" or "security@," which we've designated as "Other."

As individual teams, Akamai's finance and human resources departments are the most targeted groups, a fact that is partially masked when aggregating data as we've done in Figure 4. As in many companies, the finance and human resources teams in our company are much smaller than the teams they support. Although the count of attempts may be lower, phishing attacks per person are much higher in the teams with access to sensitive information.

We did find it somewhat surprising to see the volume of phishing attempts (12%) aimed at our executives, which includes the legal department. But given the value of compromising the account of a CEO or lawyer, it's no wonder there's so much attention being paid to their accounts.

Although phishing-related DNS requests accounted for 66% of the threats blocked by ETP, they are not the only effective use of the technology.

Nearly a quarter, 24%, of the blocked requests came from malware. This could be a link clicked on a site or a malicious document, but without manual investigation it is usually impossible to know. A further 9% of traffic came from requests to the command and control (C2) infrastructure of botnets, indicating a system that's compromised. Because the DNS request was blocked, an infected system had no method of communicating with the C2, allowing time for cleanup. The last 1% of traffic is a catch-all and might make for another research project in the future.

It's important for security vendors to use their own tools to protect themselves. It shows confidence in their technology, as well as providing insight into the challenges our customers face. But no technology is perfect, which is why Akamai builds layers of overlapping controls to protect itself – which makes us exactly like every other large organization.



Fig. 4 - Engineering, finance and HR, and marketing and sales teams receive significantly more phishing attempts than other departments within Akamai

Phishing Software Development Lifecycle

Phishing attacks follow a process and methodology that is similar to a software development lifecycle (SDLC). Starting with kit development, the cycle moves on to attack selection, propagation, and sales. This circle includes software checks, adjustments to targets or landing methods, data-collection tuning, and sales channel development. The aim is to deliver (at the lowest cost in terms of resources and finances) robust phishing kits that are secure, easily deployed, and constantly updated.

Development

Developers will design phishing kits as a near-perfect clone of the target's website. You commonly see this with kits targeting Apple, Microsoft, Amazon, PayPal, banks of all sizes, and retail operations.

The phishing kits are created to be almost perfect, because the criminals are hoping a visual inspection is all the victim will do before entering credentials, providing personal information, or downloading a file.

Kit development goes beyond basic looks, and there has been a shift in recent years toward security. Criminals don't have any issues with stealing from other criminals, so phishing kit developers have turned to sophisticated licensing schemes (as seen in Figure 5), as well as code obfuscation to keep their kits protected. These schemes don't always work, so their code might wind up as a jumbled mash-up of recycled code in someone else's kit.



Fig. 5 - A phishing kit targeting Netflix users requires a licensing scheme to prevent piracy

This leads to an interesting problem, as Akamai [wrote about earlier in mid-2019](#), wherein such recycled code leads to software vulnerabilities. These vulnerabilities were introduced after the code was copied in some cases, or existed in the original kit because of flimsy coding practices or reliance on outdated open source code. For domain administrators, a phishing kit with vulnerabilities adds additional problems to an already bad situation.

The security side of phishing kit development also focuses on evasion techniques. There are several techniques, and many kits layer them in the hope of remaining hidden for longer periods.

Akamai [has previously discussed evasion techniques](#), but some of the more common elements include geographic-based limiters, where only victims from a certain area are allowed to access the kit, and real-time text obfuscation, which prevents crawlers from finding the landing page.

Some kits also filter based on USER-AGENT and DNS resolution, looking to exclude visitors using Tor, for example, those coming from addresses associated with security vendors, or those coming from other large Internet companies such as Google or Amazon.

Many times, these kits also include an IP-based blacklist that will drop connections if they come from one of thousands of pre-configured IP sets that are known to belong to security organizations (Kaspersky, Microsoft, Symantec, Trend Micro), Internet companies (Google, Amazon, Netcraft), or universities.

Other evasions include random URLs, randomly generated subdomains, and randomly generated URI data, which adds an “official” looking random string to the URL. Often these random strings mask the domain from view on smaller screens – something that is frequently used for mobile device targeting.

Phishing kit development also includes the use of metrics, including Google Analytics. Akamai researchers recently discovered several hundred phishing websites that were using analytic tracking. While some of those sites were using the original target’s analytic ID (because the criminal copied it when they copied the target’s source code out of the browser), others were using custom unique identifiers. The custom IDs were designed to track the victim as they navigated the domain, in addition to all of the other normal analytical data that’s collected in these programs.

Attack Selection

Spear phishing attacks have a single, direct target. Generic phishing attacks are more open, and usually follow news cycles, holidays, or notable events. In the United States, Canada, and the United Kingdom, tax season is a popular time for phishing activity, but criminals aren’t forced to limit themselves to a certain time of year for their campaigns.

Targeting taxpayers is typically done for financial reasons, but there is also a good deal of personal data collection in these attacks. As such, tax-based phishing campaigns can run all year long. In early August 2019, long after the typical tax season had ended for most of the United States, Akamai discovered a phishing campaign targeting the Internal Revenue Service (IRS). Spread across 168 domains, by August 24, more than 4,000 people had attempted to visit one of the domains. As this report is being written, the campaign is still active.

The collection of personal data and financial information is why retailers are also a big target. During the Christmas holidays in the United States, for example, phishing kits for banks, online retail outlets, and even popular consumer brands like Apple, are developed and sold for a premium.

Some phishing attacks have been known to target sites related to popular vacation destinations, particularly during peak summer months for more impact. Akamai discussed this attack type in [a blog post published in August 2019](#), where we explored a phishing campaign that targeted amusement parks.

Propagation

Propagating a phishing attack, especially if it is a generic one, requires getting as many eyes on the lure as possible. In order to do this, criminals turn to customized email, social media, and SMS scripts. Sometimes a phishing attack will work better, and spread more quickly across social media, so criminals will use hashtags and other conversation points to inject their messages into the stream. The previously mentioned phishing attack against amusement park sites, for example, propagated almost entirely across social media.

Sales

There are two aspects to phishing sales. One deals with the compromised data, the other deals with the phishing kits themselves.

For the most part, the sale of compromised data happens in bulk, where combinations of usernames and passwords sell for pennies or less. Earlier this year, 50,000 usernames and passwords were selling for about \$5.

Other aspects of compromised data could be financial records (including W-2s, brokerage accounts, banking accounts, etc.) or other services such as streaming media, restaurant accounts, travel accounts (including services like Uber, Lyft, and airlines), retail rewards accounts, and more. Each of these are packaged and sold, sometimes as individual units. Other times, they can be sold in bulk based on geographic location.

Collections of PII, including everything needed to create and forge new identity documents or open an online account, are usually sold in complete sets and are available based on location, credit score, and personal net worth, just to name a few examples.

For phishing kit sales, things are a little more complex. Many kit developers operate phishing as a service (PaaS) businesses, which are usually built around an admin panel that contains a number of functions (evasion, mail and mobile messaging scripts, analytics, etc.) and additional services, such as updates, form letters, and more.

How to Protect Yourself

Retailers and enterprise vendors know their brands will be targeted by criminals running phishing campaigns, and they have dedicated serious resources to fighting them. Here's how some top targets of phishing schemes address the issue on their websites:

- [Target](#)
- [Google](#)
- [Amazon](#)
- [PayPal](#)
- [Apple](#)
- [Netflix](#)
- [Lyft](#)
- [Uber](#)
- [Microsoft](#)
- [Walmart](#)
- [JCPenney](#)
- [Macy's](#)

Phishing as a Service

To give you an idea of what PaaS looks like, consider the following example.

In Figure 6, a storefront operated by a well-known phishing kit developer offers three types of advanced kits for \$99, with additional mailer services available. The developer advertises on social media, and their storefront is public. The low prices and top-tier brand targets are attractive, creating a low bar for entry into the phishing market for criminals looking to set up shop.

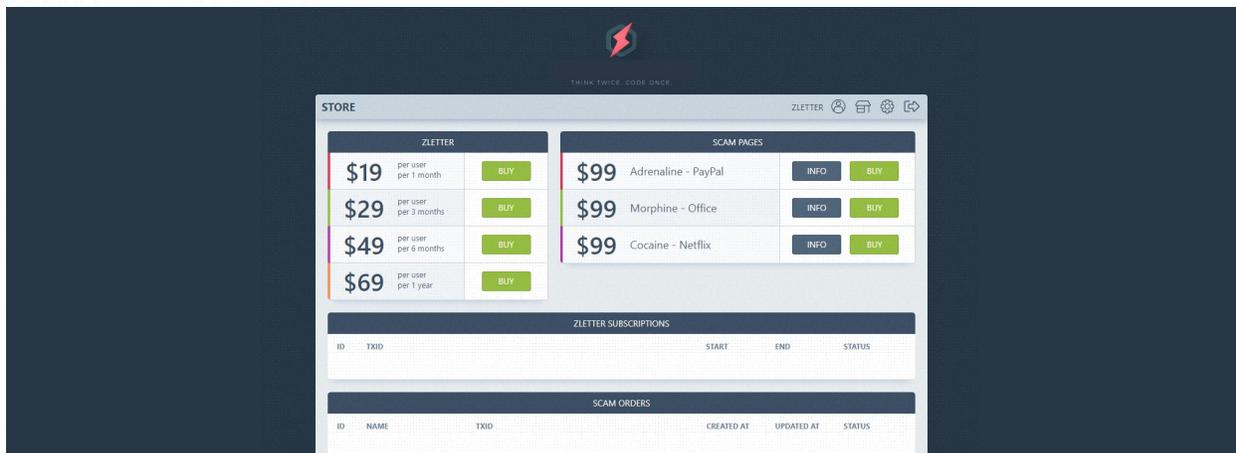


Fig. 6 - The phishing kits available in this shop have a low price point, creating a low bar of entry to criminals looking to start a new phishing campaign

In addition to the kits, this store also offers a mailing service that is priced in tiers, depending on how many months the user would like to pre-purchase. In Figure 7, the admin panel for this mailer service offers a number of options, including priority settings, random message IDs, and three types of encryption, as well as sender email and name randomization.

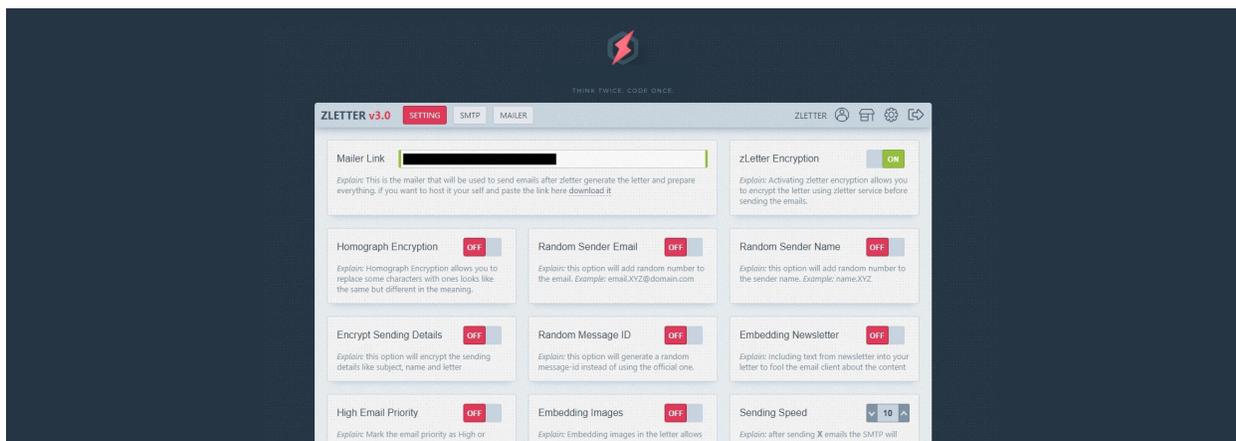


Fig. 7 - Email phishing has gotten harder over the years, requiring criminals to develop more techniques to assist in getting their messages delivered

All of the kits sold by this developer have pre-built security features and evasion techniques, but the selling point is the type of data that can be collected, and the promise of constant updates.

The kit in Figure 8 features several data collection points, free updates, and multi-language support. The kit is responsive, too, meaning it will display perfectly on a PC as well as a mobile device. For those who want a demo, the video walks potential customers through the kit and its features.

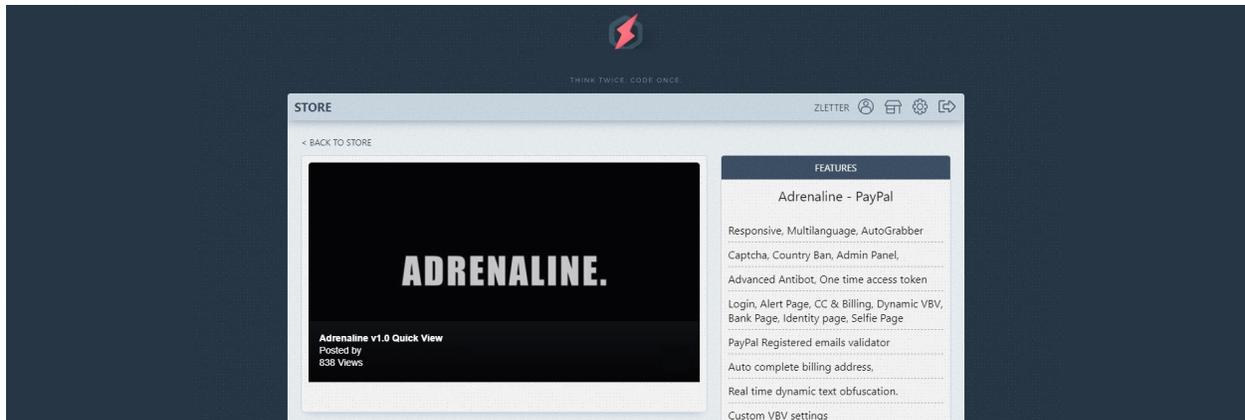


Fig. 8 - The kits sold by this developer, such as the one seen here, contain advanced features and data collection options

Another feature this developer packages into its administration panel is basic statistics. In Figure 9, some of the basic statistics are shown for a kit that was in development at the time the feature was being promoted. However, it is common for kits being sold to include stats of some kind, including functions that track victims in real time.



Fig. 9 - Basic stats pages are starting to become common in phishing kits, while some kits are leaning toward using analytic platforms such as those offered by Google

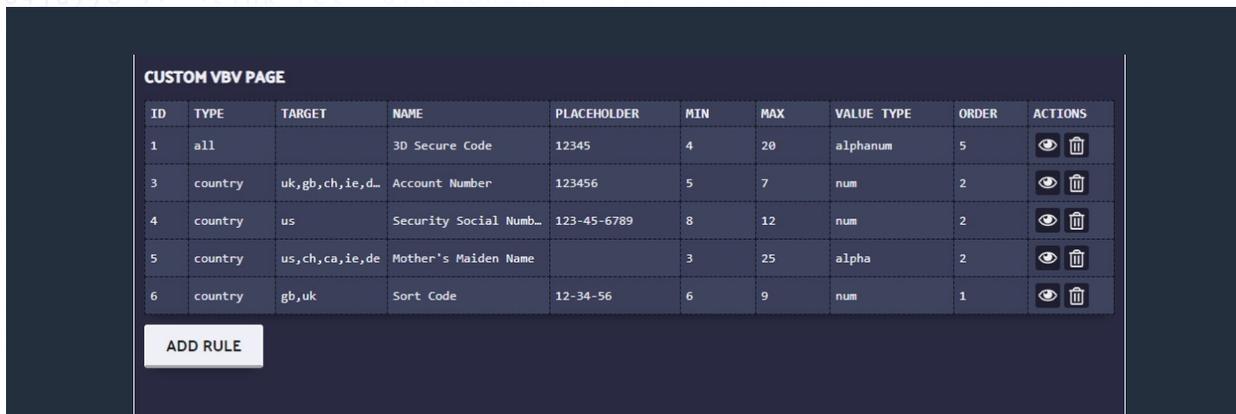


Fig. 10 - Custom data-collection templates are an uncommon feature in phishing kits, but some of the more advanced developers include them in order to stand out in the market

Customization is a feature offered by this developer. In Figure 10, one of the administrative panels enables the user to customize data collection by victim location and allows attackers to set a template to ensure that the proper type of information is collected before the victim can move on to other parts of the kit.

This developer is just one of the many who design phishing kits and developed a business model to match it. Many phishing kit developers have legitimate jobs in the tech industry, but choose to develop scam pages and phishing kits as a way to hone their skills and earn a side income.

Looking Forward



Phishing Isn't Going Away

As long as humans exist, criminals will seek to take advantage of human nature. The retail industry is just as much of a target as any other because it has a wealth of personal and financial information.

Awareness training works, and people can be trained to spot and report basic, generic phishing attacks – but this isn't a silver bullet. In fact, criminals have adapted to many basic awareness training models, which is where the boom in BEC attacks came from. By targeting the natural workflow of a victim, criminals gained a distinct advantage, which resulted in billions of dollars in fraudulent wire transactions, and millions of compromised W-2 records.

It's said that a good defense requires a good offense. So phishing simulations, combined with solid endpoint protection, can help an organization keep ahead of the phishing game and lower the odds of experiencing a disastrous incident.

However, phishing simulations need to be customized and tailored to the individual or business unit. Examples include a phishing simulation that targets HR employees by spoofing resume submissions for a recent job posting, and a simulation that spoofs lead generation responses to sales employees after a recent event.

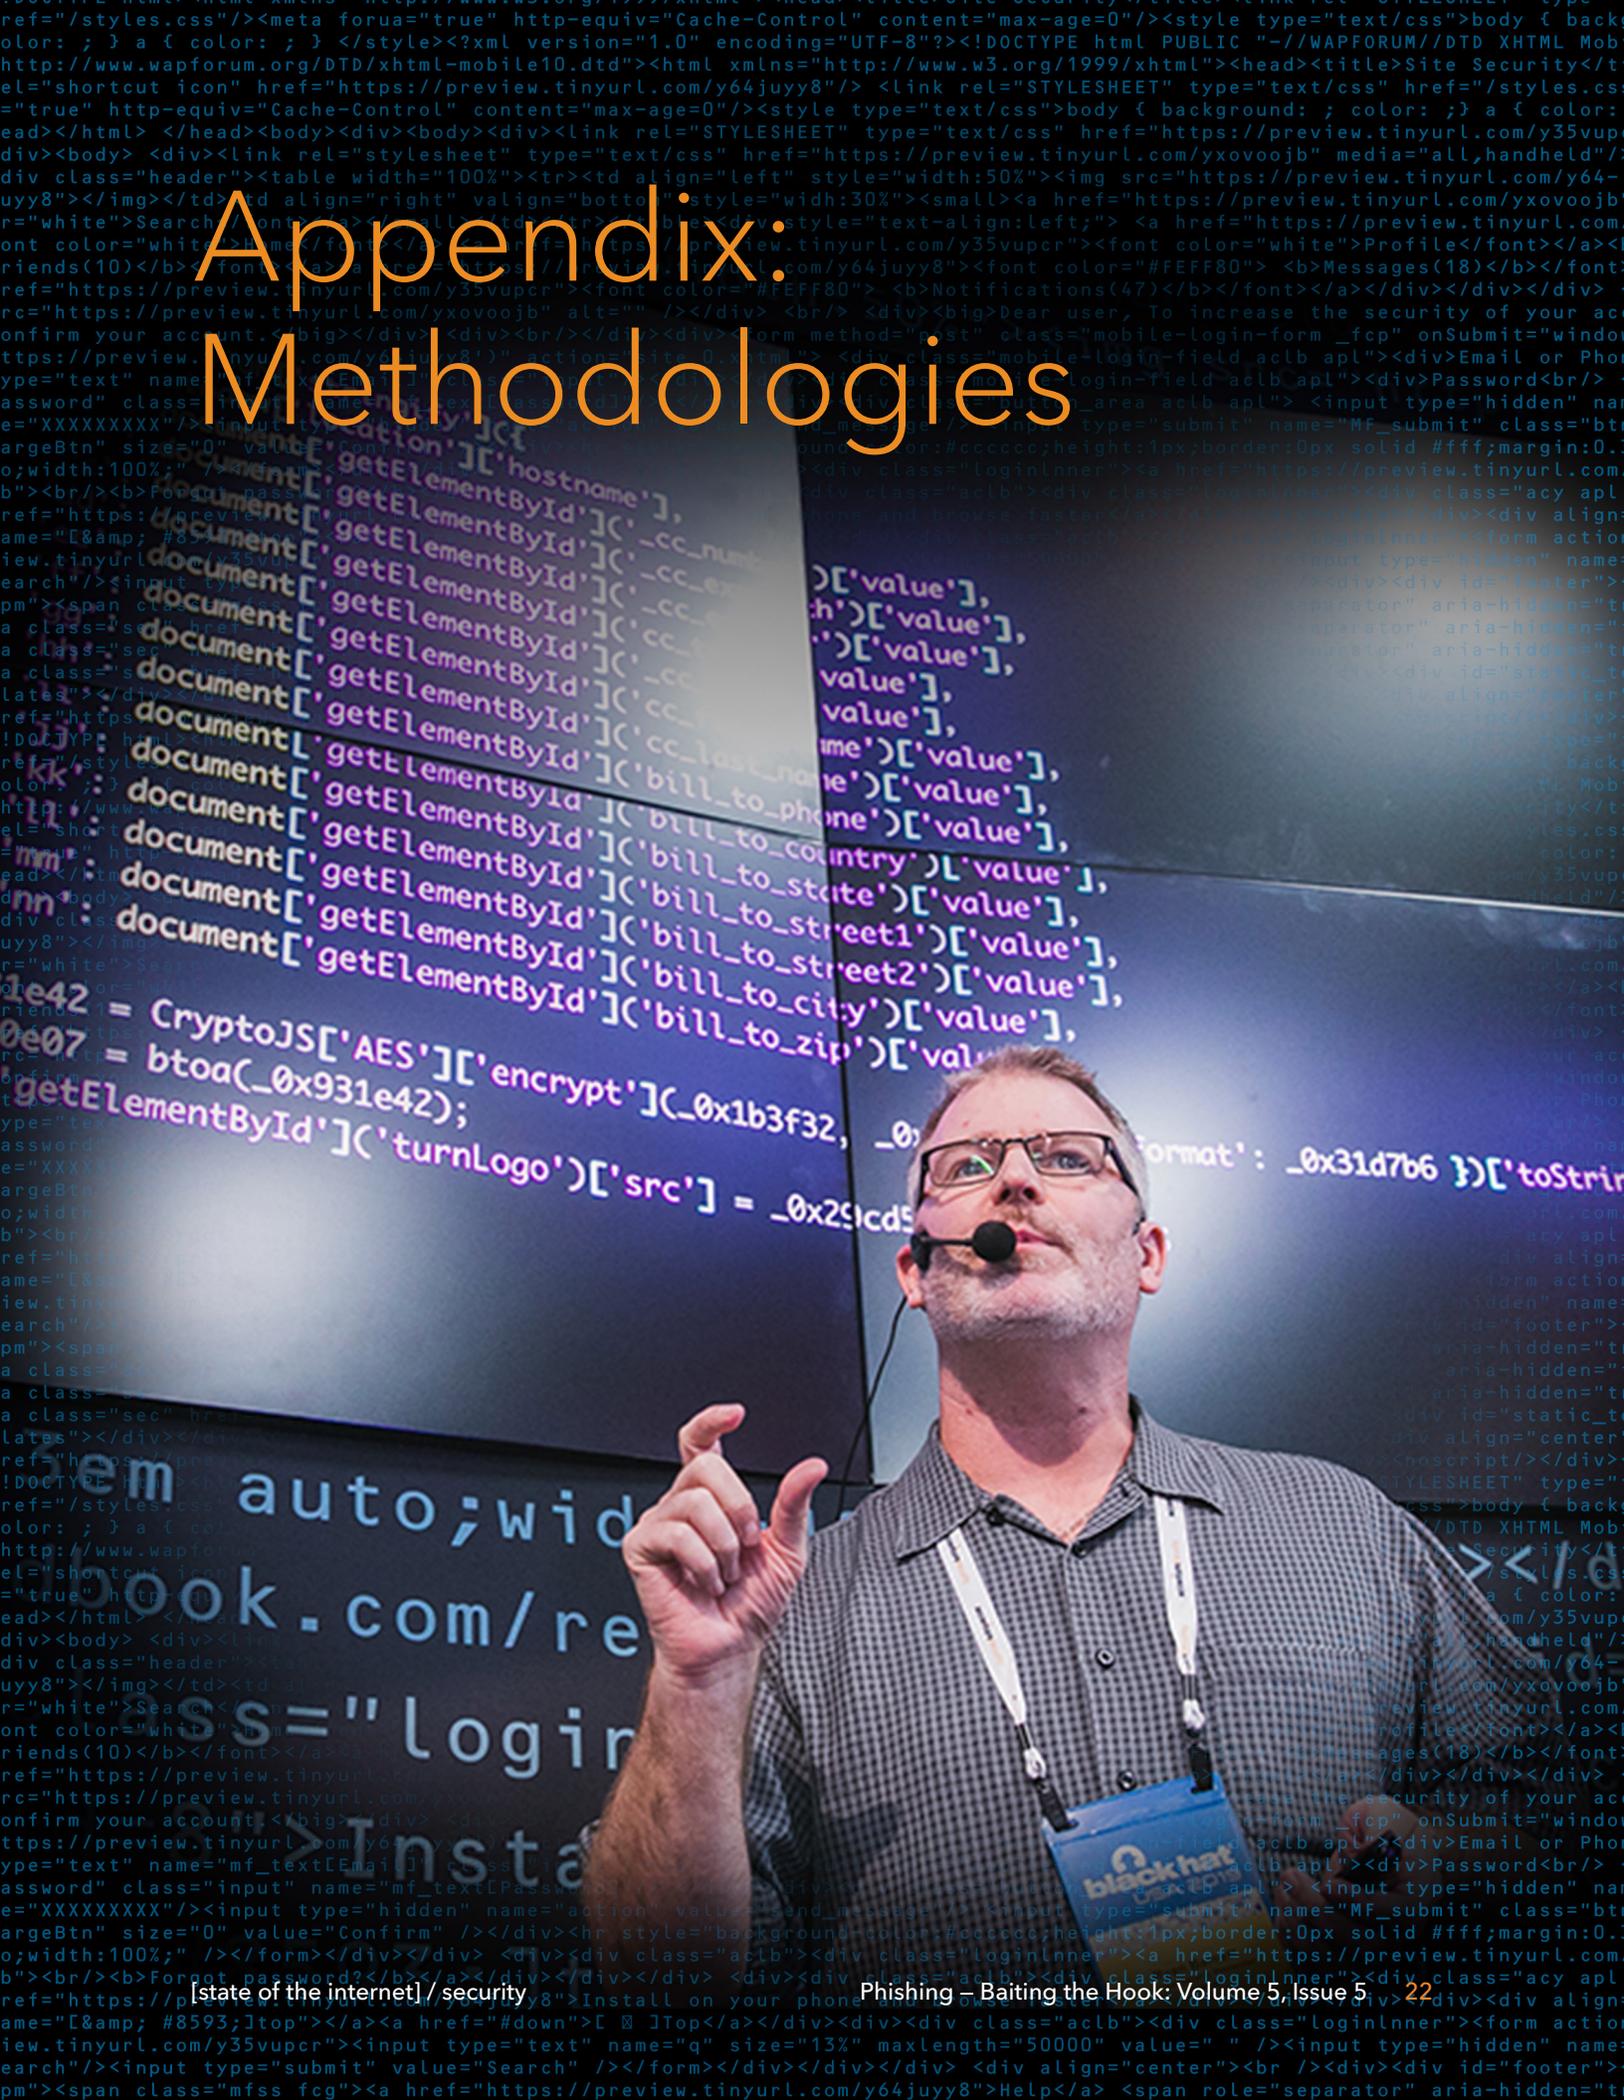
Even generic phishing simulations can be improved. Instead of spoofing a random “click here to get a prize” from some no-name company, a criminal might spoof a prize from a restaurant or retailer that employees are known to frequent or that the organization has a relationship with. Does HR have an employee perks portal? Scammers might use the knowledge of this portal as the base for an attack, where employees are offered extra perks in exchange for registration on the “new” portal using their network credentials. This same situation can also be used for awareness campaigns.

“ *Phishing simulations, combined with solid endpoint protection, can help an organization keep ahead of the phishing game.* ”

Some phishing attacks are loud and easy to spot, but lately, that hasn't been the norm. As phishing expands beyond email, new attacks can come from people and places that are known and trusted by the victim. This makes it infinitely harder to track and stop. Not impossible, mind you, just more difficult.

Phishing is not a “one-size-fits-all” attack – no two are alike. The good ones are subtle, and even the experts can be fooled by them.

Appendix: Methodologies



General Notes

The team that creates the State of the Internet / Security report does their best to make our data as clear and accurate as possible. This data and its analysis undergo multiple rounds of review prior to publication. The data used in this issue represents the efforts of multiple teams within the organization.

Using Akamai to Defend Akamai

The team responsible for securing Akamai's internal systems, Enterprise Security, provided the expertise and data for this section regarding the tools we use to protect ourselves. Enterprise Threat Protector (ETP) proactively identifies, blocks, and mitigates targeted threats such as malware, ransomware, phishing, DNS data exfiltration, and advanced zero-day attacks. ETP uses multiple layers of protection (DNS, URL, and inline payload analysis) to deliver optimal security and reduced complexity, without impacting performance. The data in this section is the result of Akamai's internal use of ETP, as well as other layers of defense.

Phishing Kit Breakdown

The data in this section is sourced from Akamai's own internal data, including traffic and proprietary phishing research over a period of 262 days. Additional information, including the details provided by Steve Ragan in Phishing as a Service, was sourced from public advertisements on social media, and darknet research.

This section also includes global DNS traffic as observed by Akamai, and internally developed tools for anomaly detection and scoring of domains and TLDs. The DNS analysis was captured from June 15 until August 14, 2019. Data was provided by the ETP and Nominum teams.

Footnotes

1. https://iang.org/papers/market_for_silver_bullets.html
2. https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=200120020SB1386
3. <http://veriscommunity.net/>
4. <https://www.darkreading.com/threat-intelligence/9-years-after-from-operation-aurora-to-zero-trust/a/d-id/1333901>

More State of the Internet / Security – akamai.com/soti

More Akamai Threat Research – akamai.com/threatresearch

Credits

State of the Internet / Security Contributors

Eric Kloster

Engineering Director –

Using Akamai to Defend Akamai

Lorenz Glaser

Senior II Security Engineer –

Using Akamai to Defend Akamai

Or Katz

Security Researcher, Principal Lead –

Phishing Software Development Lifecycle,

Phishing Kit Breakdown

Lydia LaSeur

Data Scientist –

Phishing Kit Breakdown

Paul O’Leary

Principal Data Scientist, Threat Intelligence Engineering –

Phishing Kit Breakdown

Editorial Staff

Martin McKeay

Editorial Director

Steve Ragan

Senior Technical Writer, Editor

Marketing

Georgina Morales Hampe

Project Management, Creative



Akamai secures and delivers digital experiences for the world’s largest companies. Akamai’s intelligent edge platform surrounds everything, from the enterprise to the cloud, so customers and their businesses can be fast, smart, and secure.

Top brands globally rely on Akamai to help them realize competitive advantage through agile solutions that extend the power of their multi-cloud architectures. Akamai keeps decisions, apps, and experiences closer to users than anyone – and attacks and threats far away. Akamai’s portfolio of edge security, web and mobile performance, enterprise access, and video delivery solutions is supported by unmatched customer service, analytics, and 24/7/365 monitoring. To learn why the world’s top brands trust Akamai, visit www.akamai.com, blogs.akamai.com, or [@Akamai](https://twitter.com/Akamai) on Twitter. You can find our global contact information at www.akamai.com/locations. Published 10/19.

[state of the internet] / security

Rohit Murthy

Enterprise Security Operations Manager –

Using Akamai to Defend Akamai

Keith Hillis

Enterprise IT Risk and Security, Senior Director –

Using Akamai to Defend Akamai

Steve Ragan

Senior Technical Writer –

Phishing in a Nutshell, Phishing Kit Breakdown, Phishing

Software Development Lifecycle, Phishing as a Service

Omri Hering

Senior Data Analyst –

Using Akamai to Defend Akamai

Amanda Fakhreddine

Senior Technical Writer, Managing Editor

Lydia LaSeur

Data Scientist

Murali Venukumar

Program Management, Marketing

Phishing – Baiting the Hook: Volume 5, Issue 5 24