Volume 5, Issue 4

Financial Services Attack Economy

Table of Contents

01 | Letter from the Editor

The attacks and tools being used against financial services organizations are part of a complex ecosystem. This may seem like an obvious statement at first glance. When you read most industry reports, they tend to focus on a single aspect of the whole ecosystem, as if it were unrelated to other types of attacks. To be honest, we've been just as guilty of having a myopic view of attack traffic as anyone else.

We wanted to do something different for this report. Instead of looking at a single type of attack, we stepped back to look at attacks against banks, credit unions, trading companies, and other organizations that make up financial services as a whole. Most defenders only see a very small segment of the overall traffic, whether they're the target or the vendor supplying defensive tools. The breadth of Akamai's products and our visibility into a significant portion of Internet traffic allows us to research multiple stages of the attack economy.

There's no one organization that can detect and track all aspects of the attacks we face today. But combining intelligence from different types of attack data with broad experience allows talented analysts to understand the attacks they do see, and how they are related to the environment as a whole.

Rarely does a week go by without news of a data breach, and, inevitably, compromised accounts. Many of the usernames and passwords gained from these attacks end up in large "dumps" that are bought, sold, and traded to fuel new attempts to compromise sites.

This edition of the State of the Internet / Security report uses passwords as the starting point of our analysis and follows the data through multiple aspects of the criminal economy. Bank, credit union, and other financial institution accounts represent the pinnacle target for attackers, but the ultimate goal is, as always, the transfer of money from the target to the pockets of criminals. Whether it's phishing, credential abuse, Distributed Denial of Service (DDoS), or another type of tool, money is the goal for the vast majority of attacks, and the entire ecosystem is built around that.

Every organization should be paying attention to the attacks targeting financial services systems. The tactics, tools, and procedures being used against banking services will expand to other types of targets if they show signs of being successful. Like many of the warnings we see in security, it's not a matter of *if* your data and systems are going to be targeted, it's a matter of *when*. The ecosystem doesn't stop with financial services.

TL;DR

- **50%** of all the unique organizations impersonated by tracked phishing domains were from the financial services sector, according to Akamai's records
- We observed that more than **6%** of global malicious login attempts targeted the financial services sector
- **94%** of the attacks against the financial services sector came from one of four methods: SQL Injection (SQLi), Local File Inclusion (LFI), Cross-Site Scripting (XSS), and OGNL Java Injection (which accounted for more than 8 million attempts during this reporting period), based on Akamai's calculations

02 | Recent SYN-ACK Reflection Attacks

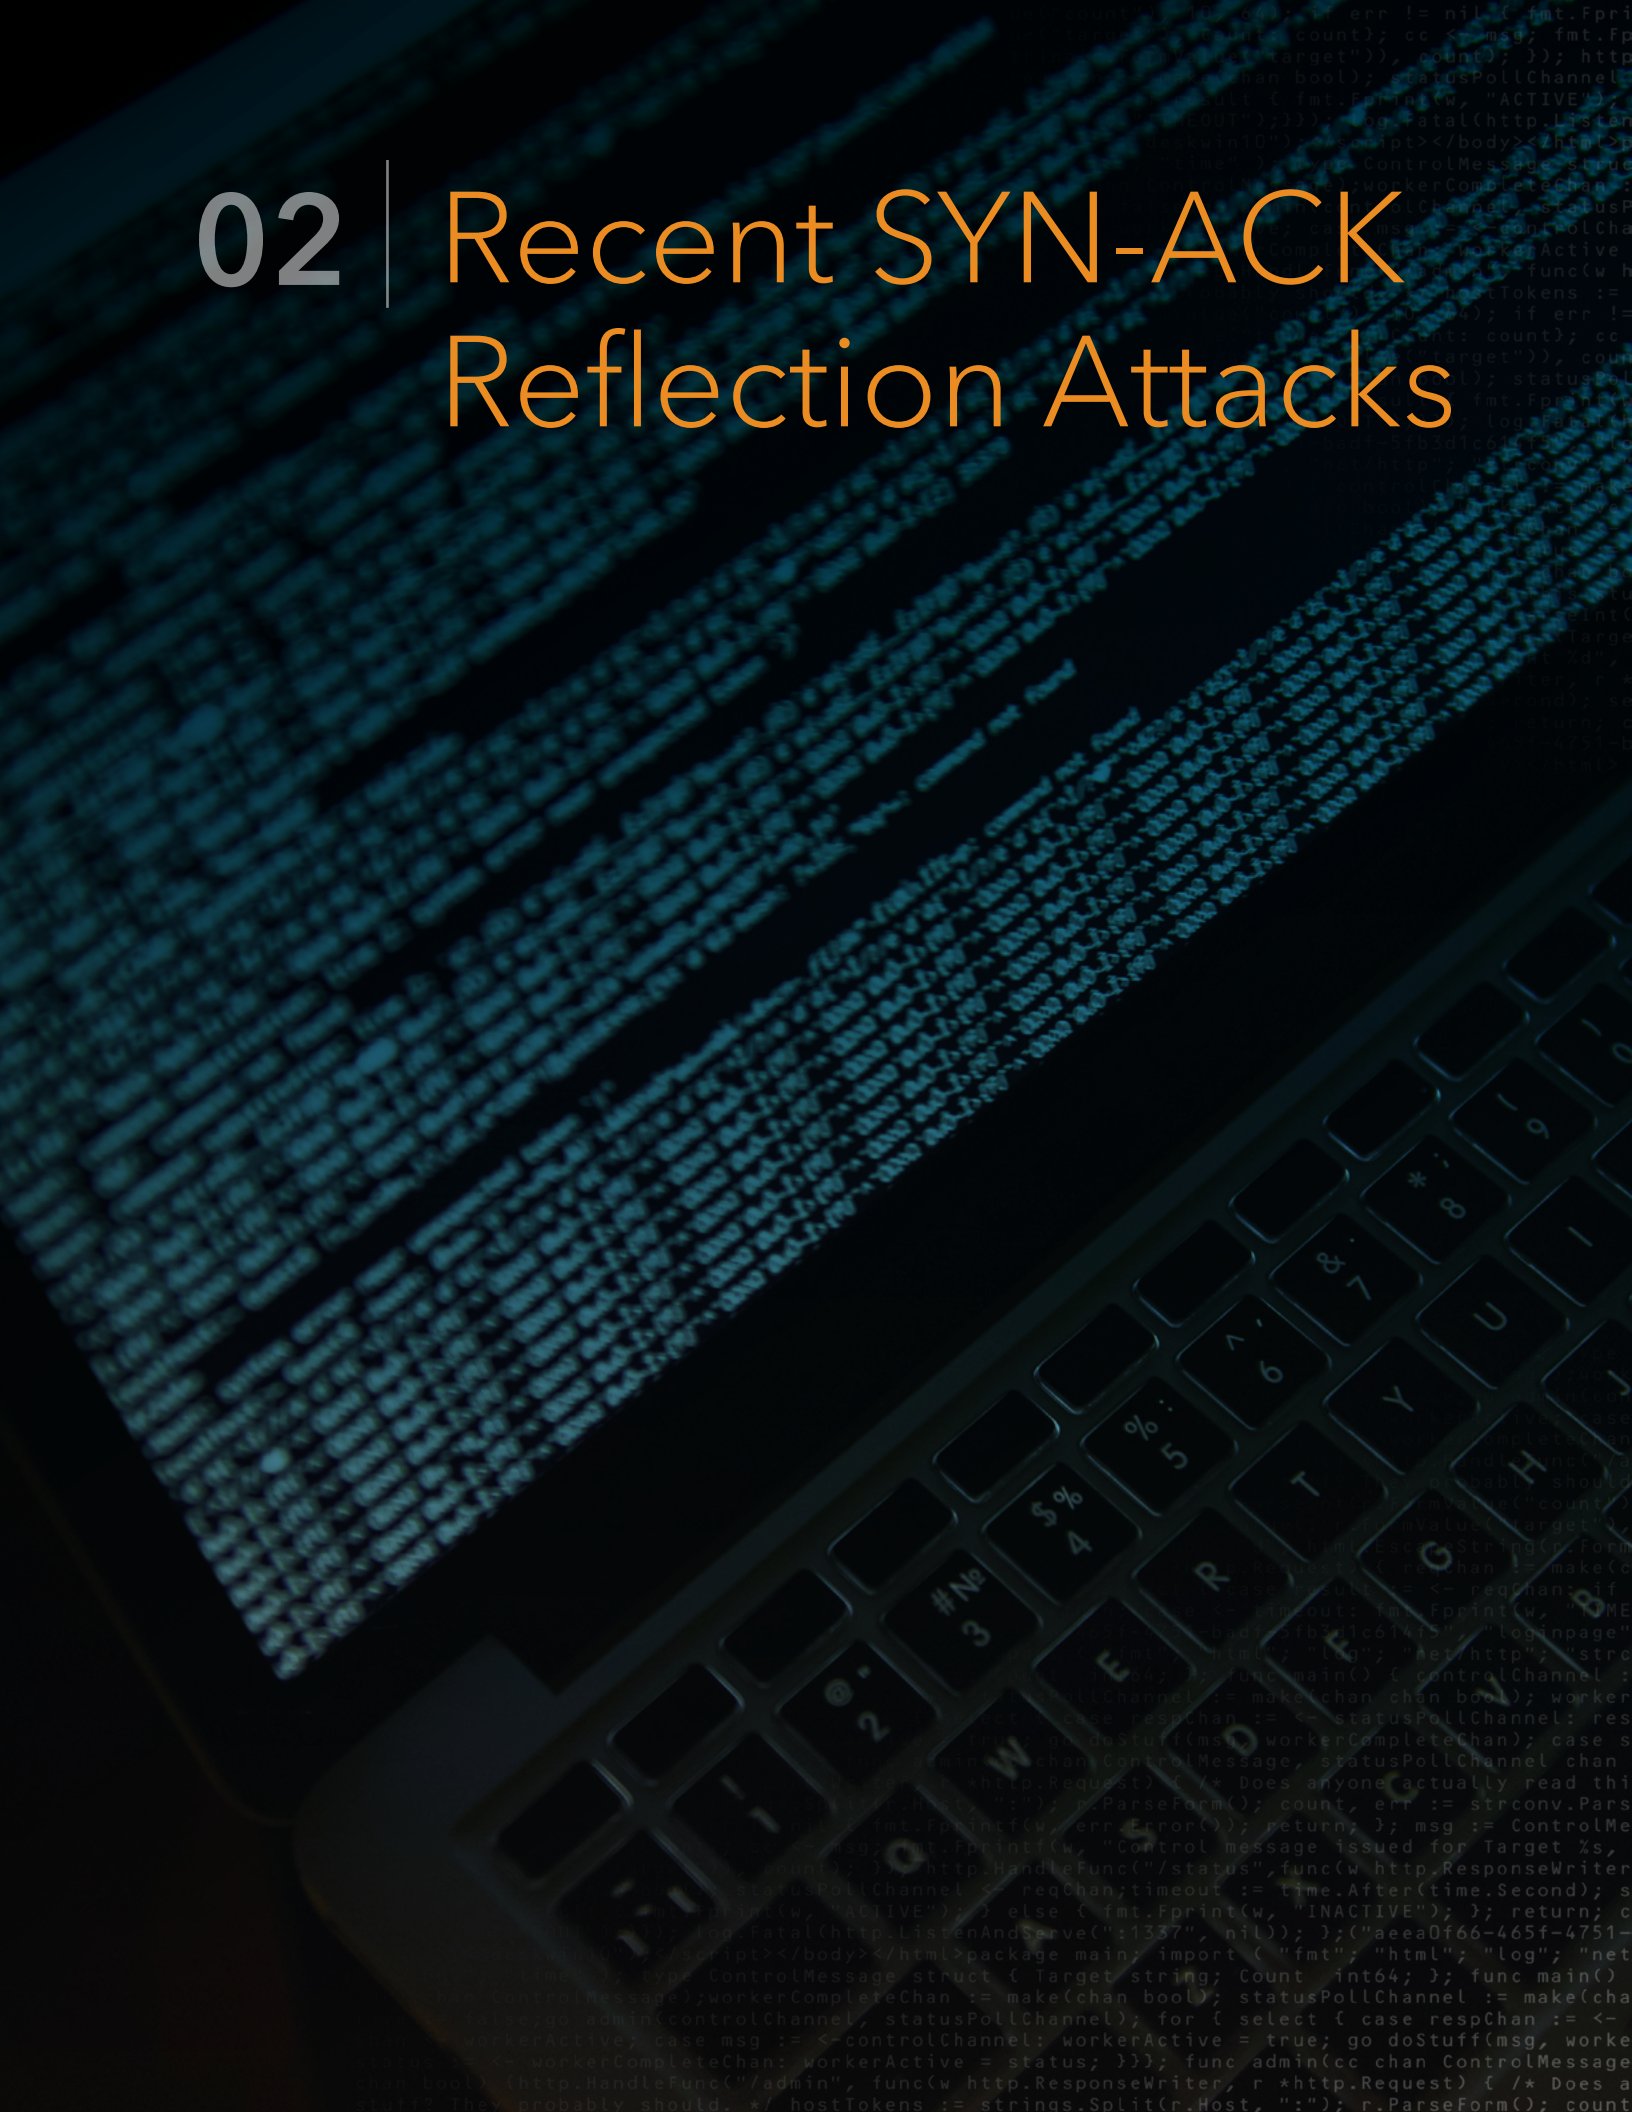

In March 2019, several financial services organizations started seeing DDoS attacks using TCP SYN-ACK packets to flood their data centers. This type of attack is not commonly used by attackers because of its limited impact on the target. To put it bluntly, this is the type of attack someone with minimal technical knowledge and understanding of networking would use.

But further research into these attacks suggests they're more than they seem on the surface. Because this style of attack generally has a minimal impact on the target, the inclination is often to ignore the traffic. Two factors made these attacks different from prior SYN-ACK floods: the number of targets being affected and the secondary effects stemming from this traffic.

Because TCP SYN-ACK floods are relatively uncommon, these attacks were easier to tie together than other styles of attacks Akamai commonly sees. As researchers within Akamai began discussing the traffic with peers, both internally and externally, it became clear that multiple organizations were seeing similar packets on their networks. The majority of the attacks were quickly discovered to target financial services organizations.

What were harder to determine were the secondary effects of these attacks, which might have been the true intent of using TCP SYN-ACK as the primary attack vector. The original SYN packet used to cause the reflector to send a SYN-ACK was seen by the reflector as a SYN flood attack and caused the reflector to tag the spoofed IP address as a malicious actor. This created a secondary effect of blacklisting financial services IP addresses, and creating further issues for defenders.

One theory suggests the effort was to hurt the reputation of the tools, and the tool makers, by maliciously causing financial services organizations to be falsely identified as bad actors. The Spamhaus Project was especially affected, as angry security teams called to understand why they had suddenly been listed as malicious. As of this writing, no one has claimed responsibility for the attacks and there is no direct evidence of intent.

It looked like an amateur DDoS attack against financial institutions. Until the secondary effect came to light.

It's hard to find definitive proof that the secondary effect on the IP reputation of the targets and reputation tools was intentional. However, the theory that organizations like Spamhaus were the real target is bolstered by adaptations made by the attackers. Over time, their traffic was modified to include spoofed User Datagram Protocol (UDP) traffic, aiming to increase their impact. It is not unusual to see multi-protocol attacks, but it is rare to see an attacker change tactics during a campaign.

These attacks have had limited long-term effects and have dwindled in frequency. The initial effectiveness of these tactics was minimized quickly as defenders understood both the direct and secondary damage they created. Communication between organizations makes it unlikely this tactic will be effective against financial services in the future, but other types of businesses with less cohesive communication strategies might be the next set of targets.

04 | Overview

The financial services industry is one of the largest and most sought-after targets for attackers. This particular industry revolves around trust and security, and because of this, these elements are quite often the key focal point for criminals.

Successful attacks can quickly translate into the dissemination of massive amounts of personal information and monetary gains by the attacker. Criminals targeting the financial services industry do so by leveraging various attack surfaces, such as people, processes, applications, or systems.

Regardless of where the financial institutions are located in the world, staying aware of the forefront of new attack trends is imperative for keeping their systems and customers safe. Understanding the ecosystem of attacks, how different layers interact, and the booming economy spawned by it gives financial services organizations the edge they need.

05 | Akamai Research

Password Problems

Passwords are both a single point of failure and the key method of identification on networks around the globe, and have been for centuries. They're part of the human existence. Roman soldiers used "watchwords" as a method of securing areas, and Allied forces used countersigns on D-Day during World War II (e.g., the countersigns Flash / Thunder).

The concept of passwords that we are familiar with can be traced back to the early days of computing, with the development of the IBM 7094 and the Compatible Time-Sharing System (CTSS) operating system.

In the 1960s, Massachusetts Institute of Technology's CTSS was the first system to leverage the use of passwords as a means of authentication. It was also the first system to experience a password-related data breach. The incident was documented by the IEEE in a report chronicling the CTSS's 50th anniversary.

In the spring of 1962, an MIT researcher needed more than their allotted four hours per week on the CTSS. The researcher discovered where all of the system passwords were stored, and submitted an offline print request for that file. Early the following morning, the researcher retrieved those passwords, and continued their work using the newly acquired credentials.

We've come a long way from compromising passwords for research access to compromising passwords in order to fuel a massive criminal economy. However, passwords are still at the root of many problems faced by enterprises today, including those in the finance sector.

Over the years, organizations have relied on password policies that are too complex, such as requiring long strings of characters, numbers, and letters that are forcibly changed every so often. This leads to instances where credentials are easily guessed, or worse, easily obtained thanks to poor storage and management, because passwords were either written down or recycled across a number of domains and services. End users don't fare any better for nearly the same reasons.

Passwords are a double-edged sword, and managing identity is getting more difficult every day. There are gaps in identity policy that criminals seek to exploit, and it's in these gaps where the criminal economy thrives.

Password requirements have become too complex, allowing attackers to exploit user frustrations.

Staging Financial Attacks

One of the common attack methods we've covered across all of the State of the Internet / Security reports this year is credential stuffing. Certainly, the financial sector is not only aware of credential stuffing, but experiences these types of attacks daily. While financial organizations have become very good at dealing with them, no system is perfect.

But what are credential stuffing attacks? Where are the criminals getting their data from?

Credential stuffing attacks are a sub-vector of brute force attacks (see OWASP for more details). Essentially, it is the automated injection of breached username and password combinations against an authentication system, such as a login form or API.

Attackers use All-In-One (AIO) applications to automate credential stuffing at scale. AIO applications are easily obtained for little or no expense. For example, the popular SNIPR tool retails for about $20, while others can be found for free. *Free* doesn't always really mean *free*, however, since buyers can pay for custom configurations depending on the service or business they're targeting. Either way, the bar for entry is low for those looking to conduct credential stuffing attacks.

Having an AIO application is only one half of a credential stuffing attack. The other half requires a combination list. These lists are the key element to any credential stuffing campaign, and they can be obtained and compiled in various ways.

A screenshot of the SNIPR product page

One of the most common methods for generating combination lists is to focus on data breaches. Every day, compromised credentials are exposed on the Internet, or an unsecured database housing credentials is discovered online and secretly copied. Over the last decade, hundreds of millions of accounts have been exposed because of various data breaches, and a common theme started to emerge among the leaked records – recycled passwords.

Credential Stuffing Masterlists

In January 2019, downloads containing usernames, passwords, and email addresses were released to the web. Labeled "Collections #1-5", the releases amounted to about 1 TB of data. At the time, the media reported on the discovery of these data dumps as the "mother of all breaches," but the reality of the situation wasn't as dramatic.

In fact, the credentials contained within the releases were mostly collections of previously exposed usernames and passwords. One expert – Troy Hunt, who maintains the Have I Been Pwned? website – examined Collection #1 and determined that while there were 2.6 billion rows of data, only 1.1 billion unique email addresses and passwords were present. That was then drilled down to about 773 million unique email addresses, and 21 million unique passwords.

After the releases became public, collections of credentials – often sorted by email domain – started appearing on forums across the Internet. As time went on, these releases were scrubbed of bad data, re-sorted, and circulated again. The goal of the collection releases, as well as the scrubbing and recirculating of the data, was to target accounts that are likely to have shared or recycled passwords.

Recycled passwords are why credential stuffing attacks work. A data breach on a sports forum or video game can easily transform into compromised email and banking accounts if those credentials are recycled. Sometimes (but not always), credential stuffing attacks will leverage permutations, instead of relying on the original password list. While the compromised password might be "Scott123," a permuted list will include that, along with "scott123," "Scott321," "Scott1234," and so on.

Using permutations isn't a consistent practice among criminals and usually only happens when there is a strong focus on a specific target, whether it's an organization or an individual. If an email address or username is discovered in multiple data breaches, and there are different passwords associated with it, a criminal will attempt to use all of them, as the odds of success are higher at that point.

While the majority of the Collections #1-5 credentials were previously known, the addition of millions of passwords and hundreds of millions of new email addresses isn't something to dismiss. There has been a steady, trackable stream of credential stuffing attacks across the Internet since these releases became known to the wider public.

Credential stuffing lists, like the ones in Collections #1-5, are usually traded or given up freely to anyone looking for them in the right places. Figure 1 shows an online forum with precompiled credential stuffing lists available for download. In the image, the lists — which are created and provided by forum users — are targeting email providers, retailers, gaming, and various online services. In addition, forum members also share complete archives of records exposed during data breaches.

However, some of the data in these lists is also sold on the darknet. The credential stuffing lists being sold are sorted and focused toward a number of Internet services, such as gaming, entertainment, and financial services. The cost associated with these lists will vary depending on a number of factors, including geography, targeted service, freshness, and volume.

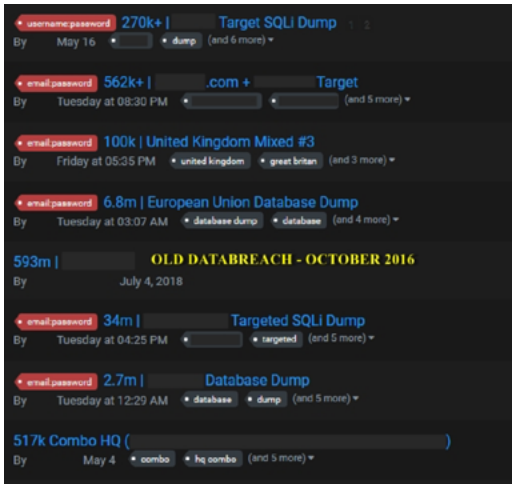

Fig. 1 – An online forum offering precompiled credential stuffing lists for download

For example, a list of 50,000 email addresses and passwords that have been sorted by mail provider and location, could go for about $5.50. However, that same list, focused toward an online financial service or local bank, could sell for triple that amount. Then again, bulk, nonverified lists can sell for pennies.

Phishing

Freshness is a commodity on the darknet when it comes to buying and selling anything, like credentials or credit card data. Data breaches are one way to get fresh records. Another way — one that often yields instant results — is phishing.

Phishing is a well-known attack vector in the security industry and in the financial space. However, despite massive amounts of effort poured into awareness campaigns, phishing still remains a top threat.

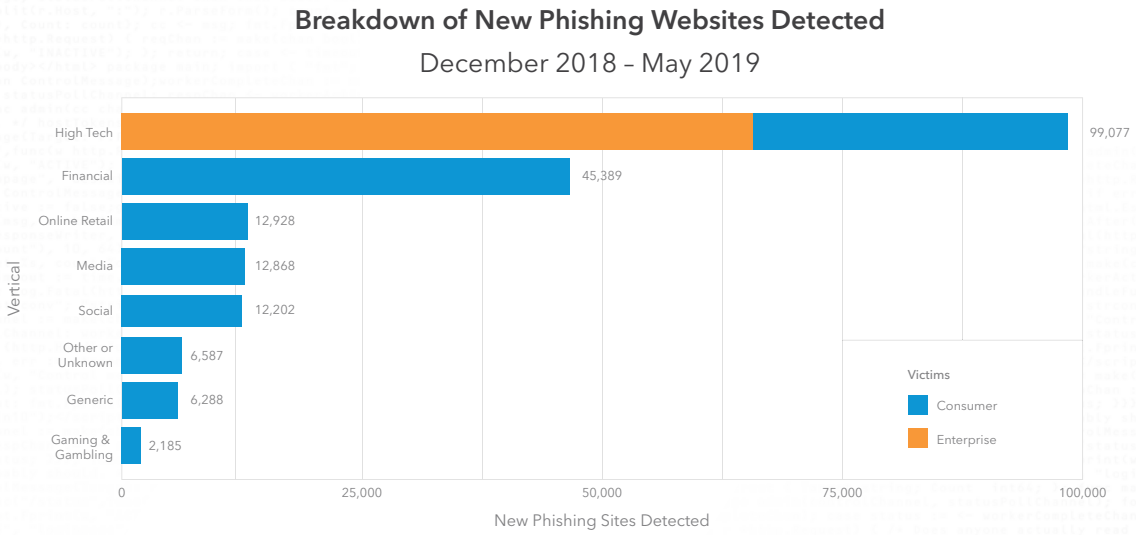

Fig. 2 – A breakdown of all new phishing websites detected between December 2, 2018, and May 4, 2019

Between December 2, 2018, and May 4, 2019, Akamai detected 197,524 phishing domains. Of those, as shown in Figure 2, 66% (130,242) targeted consumers, while the remaining 34% (67,282) targeted enterprises.

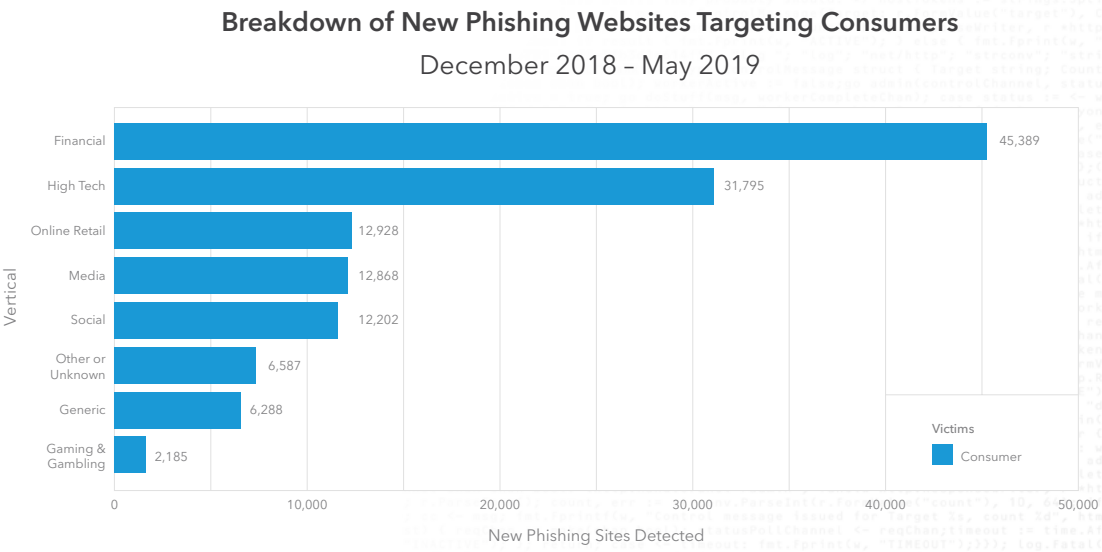

Fig. 3 – A breakdown of new phishing websites targeting consumers between December 2, 2018, and May 4, 2019

Looking at the phishing domains targeting enterprise victims, 100% of them were impersonating sites from the high tech industry. However, when only phishing domains targeting consumers were considered, financial organizations take the top spot, as illustrated in Figure 3. In fact, 50% of the unique organizations impersonated by the tracked phishing domains fell within the financial sector, according to our data.

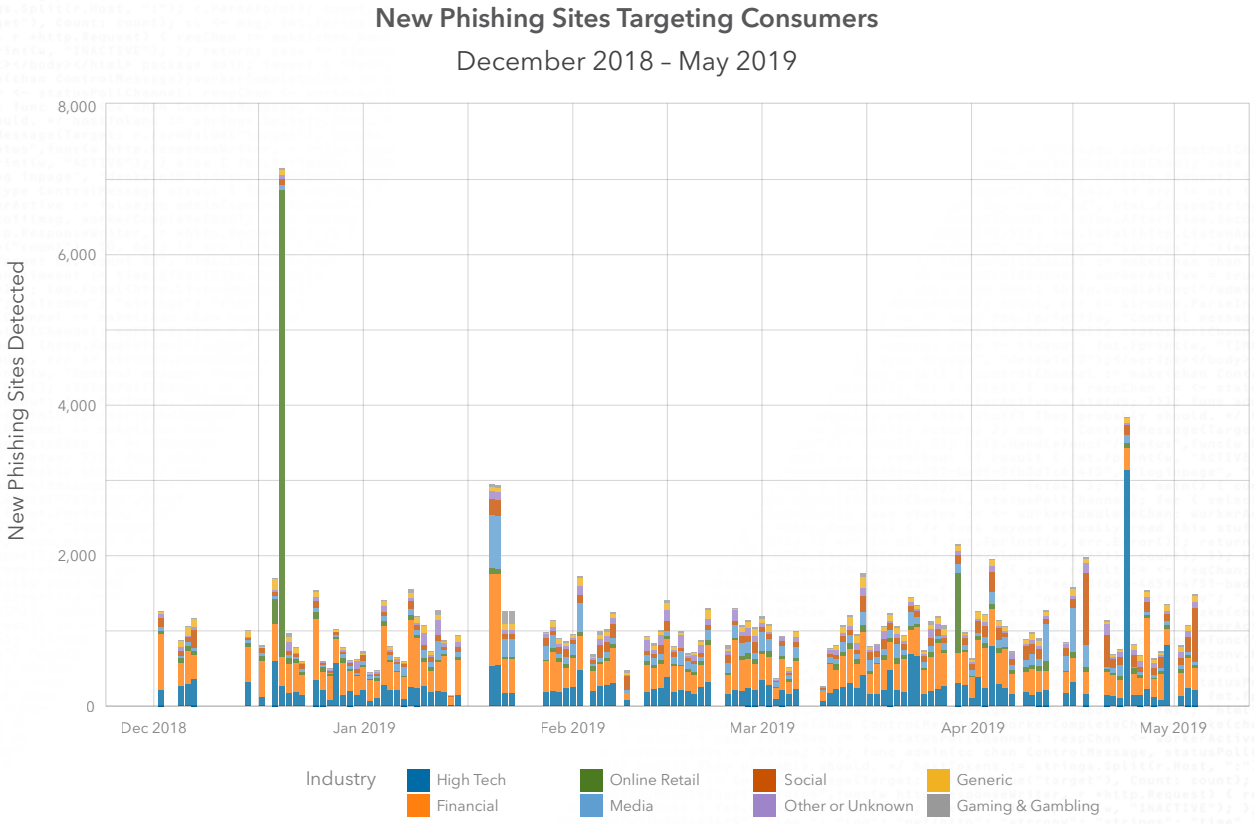

Fig. 4 – Over time, the number of phishing domains detected targeting the financial sector remains consistent

Over time, the number of new phishing domains targeting consumers has remained steady, with a noticeable spike on December 20, 2018, as shown in Figure 4. This spike, centered on the peak of holiday shopping, is not unusual since criminals often target holiday shoppers.

The impact of phishing against financial organizations should not be taken lightly. Not only are brand reputations taking a hit, but customers are placing their identities and financial security at risk each time a phishing attack is successful. Many phishing kits do more than just target usernames and passwords. Some kits also collect personal information such as passport data, driver's license data, credit card information, and more. All of these data points can then be collected and sold, or used in a number of fraud-based schemes.

Another type of phishing-related attack, which also has had a serious impact on the financial sector, is a business email compromise (BEC) attack. This type of attack targets businesses and individuals, but the overall goal is the same: Trick the victim into transferring funds directly, or release financial records that can later be leveraged in financially based fraud. According to the FBI, BEC attacks resulted in 1.2 billion dollars of losses in 2018.

Launching Financial Attacks

Once the credential stuffing lists are obtained, and the AIO applications configured, the criminal needs to actually initiate an attack. Security is a top concern since the financial services industry is so regulated. This means a simple point and shoot attack – just blasting a bank with usernames and passwords in rapid succession – isn't going to work. Most attacks against financial services take a low and slow approach.

The low and slow approach requires that the criminal uses several proxy configurations and limits their processing to a few dozen attempts at a time. While Akamai data shows that not all criminals take the low and slow approach, the ones that do are more difficult to catch.

Credential Stuffing by the Numbers

Akamai tracked 18 months of credential stuffing attacks for this report, covering November 2017 through April 2019. During this time, we observed 57,970,472,311 malicious login attempts, of which 3,547,533,230 were against financial services organizations. Overall, 6.1% of the malicious login attempts globally were targeted toward the finance sector.

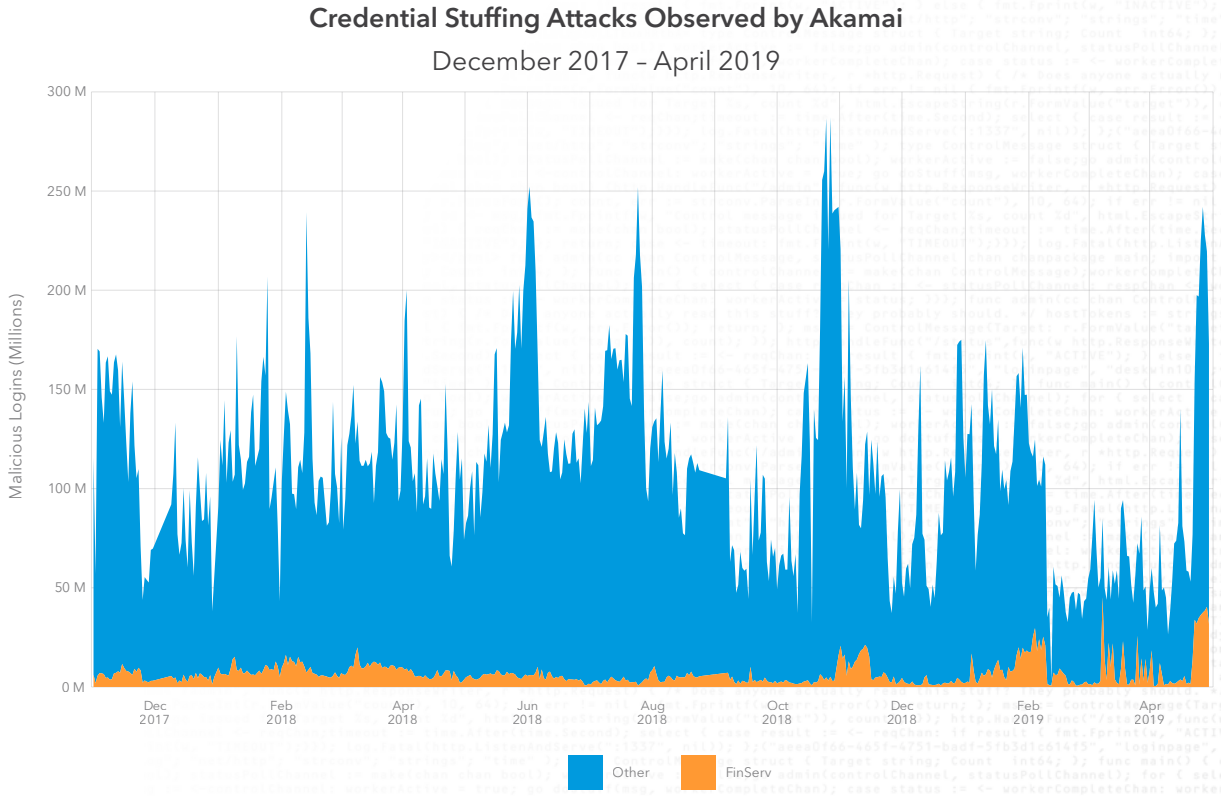

Fig. 5 – Credential stuffing attacks observed by Akamai during an 18-month reporting window

Globally, when it comes to malicious logins against financial service organizations, the United States took the top spot. The United States was then followed by China, Malaysia, Brazil, and Germany, to round out the top five.

Top 5 Countries Responsible for Credential Stuffing Attacks

November 2017 – April 2019

COUNTRY	GLOBAL RANK	MALICIOUS LOGINS	FINSERV MALICIOUS LOGINS	PCT FINSERV
United States	01	18,542,844,141	1,133,026,190	6.11%
China	05	2,077,424,958	251,719,283	12.12%
Malaysia	14	982,450,816	227,443,763	23.15%
Brazil	03	3,197,885,812	173,176,382	5.42%
Germany	12	1,357,470,150	137,863,141	10.16%

Fig. 6 – Ranking of the top 5 countries responsible for credential stuffing attacks during the reporting period ***(Top 20 list located in Appendix B: Supplemental Data)***

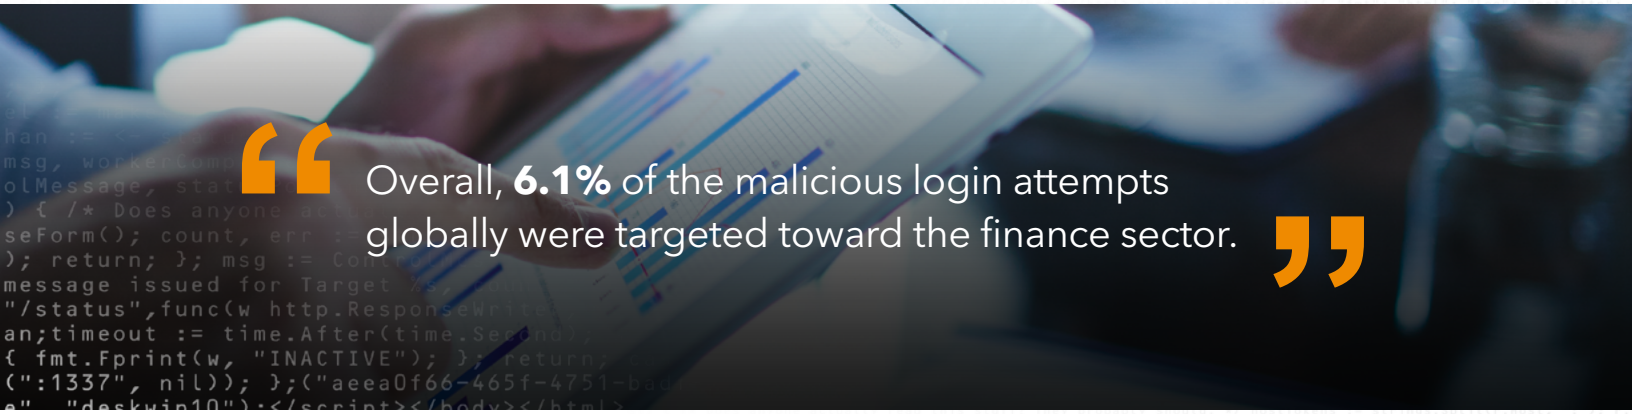

Web Attacks by the Numbers

Sometimes, criminals will attempt credential stuffing attacks side by side with distractions, such as DDoS attacks, or they'll skip the credentials and attempt to exploit applications or website vulnerabilities on the target's domain.

Looking at overall web attacks during the same 18-month period, we observed 4,460,367,847 attacks across all verticals. When drilling down to just attacks against financial services, just over 9% (411,409,583 attacks) impacted this segment.

At the same time, the financial services industry accounted for 14% of all unique targets during this period. Over time, even as the volume of web attacks increases, we can see that attacks against financial services have remained relatively stable, as seen in Figure 7.

Web Attacks Against Financial Services

November 2017 – April 2019

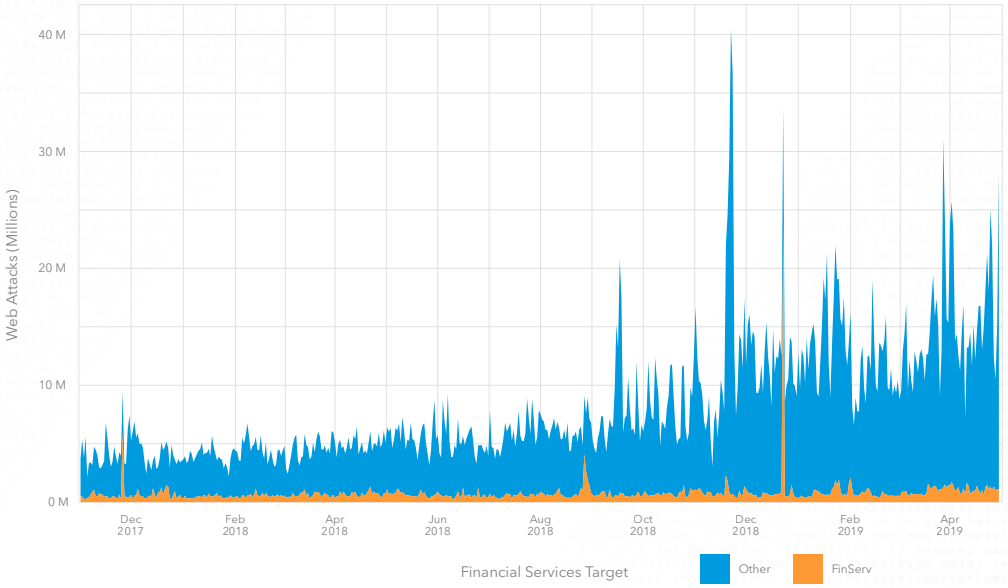

Fig. 7 – Over time, as web attacks increase, attacks against financial services web attacks remain relatively stable

Attack Types by the Numbers

Figure 8 illustrates that 94% of the attacks against the financial services sector reviewed by Akamai came from one of four methods: SQLi, LFI, XSS, and OGNL Java Injection.

Web Attacks Ranked by Vector

VECTOR	TOTAL ATTACKS	NUMBER OF FINANCIAL SERVICES ATTACKS	PERCENT OF FINANCIAL SERVICES ATTACKS
SQLi	2,967,809,139	171,305,214	41.64%
LFI	1,050,852,414	166,790,717	40.54%
XSS	188,426,146	40,682,394	9.89%
OGNL Java Injection	31,827,918	8,569,324	2.08%
PHP Injection	103,604,265	8,169,148	1.99%
RFI	69,202,119	7,701,055	1.87%
Command Injection	29,332,277	5,787,050	1.41%
Malicious File Upload	19,313,569	2,404,681	0.58%
Totals	4,460,367,847	411,409,583	100%

Fig. 8 – SQLi still remains the most common attack vector against financial services, but the inclusion of OGNL Java injections is noteworthy

OGNL Java Injection accounted for more than 8.5 million observed attempts during the reporting period, and that attack type is arguably the most sensitive of the group after SQLi. Such a high volume of OGNL Java Injections serves to remind us that attacks against Apache Struts are still a popular option for criminals targeting the financial services industry, even two years after patches were made available.

> Attacks against Apache Struts are still a popular option for criminals targeting the financial services industry.

CVE-2017-9791 was disclosed on July 7, 2017, and impacted several retail, insurance, and finance-based back-end systems, including anything that ran Struts 2.3.x (prior to Version 2.3.33) with the Struts 1 plugin. This plugin allowed developers to use existing Struts 1 Actions and ActionForms in Struts 2 applications. The vulnerability enabled remote code execution on an affected server. The exploits against this CVE have been included in Metasploit and released to the public, meaning that anyone with a bit of knowledge can easily scan for vulnerable systems and target them.

As noted in Figure 9, the vast majority of web attacks we reviewed involving the financial services segment targeted banking, followed by card services, insurances, and financial exchanges.

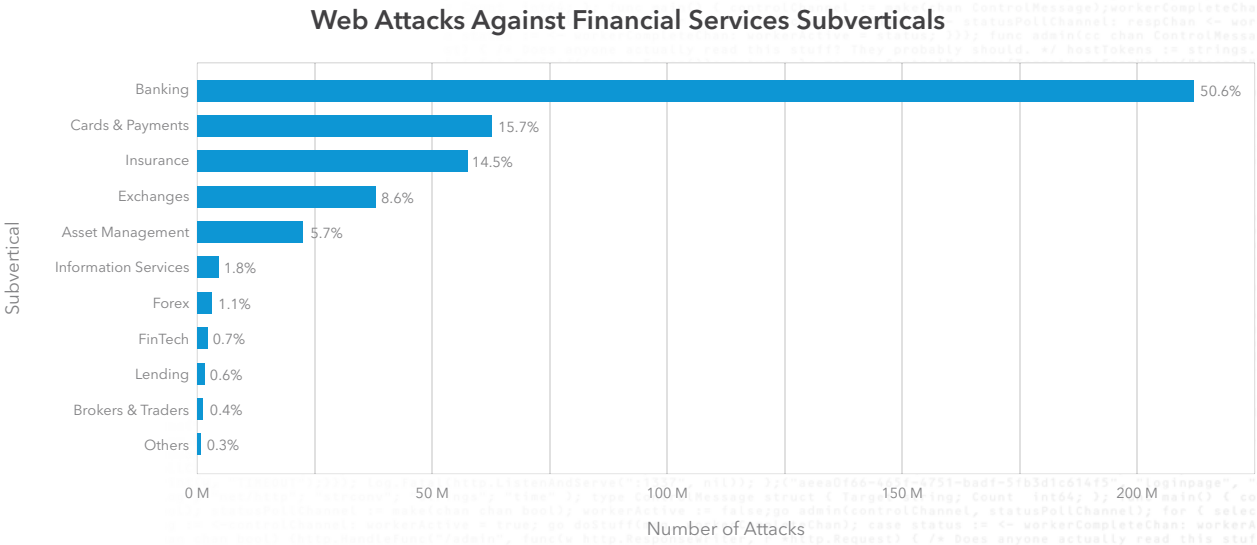

Fig. 9 – Banking was the top target for web attacks during the reporting period

A successful web attack against a financial organization can be disastrous and cost millions of dollars in cleanup and recovery efforts, not to mention the reputational impact such a data breach could have. Organizations spend considerable resources on developing secure and robust applications and public-facing web services, but things can get missed in development and quality assurance checks, which is what criminals are looking to exploit.

DDoS Attacks

DDoS attacks are a problem for any organization, but they are especially a problem for the financial services industry. A successful DDoS in the financial world could mean millions of dollars lost for each minute of downtime.

As mentioned, sometimes criminals will launch DDoS attacks as a distraction, either to conduct credential stuffing attacks or to exploit a web-based vulnerability. In the same 18-month data set, Akamai tracked more than 800 DDoS attacks against the financial services industry.

While gaming was the top target between November 2017 and April 2019 attacks in terms of attack volume with just under 9,000 attacks (Figure 10), the industry with the most unique targets was financial services. During the 18-month window, more than 40% of the unique DDoS targets were in the financial services industry (Figure 11).

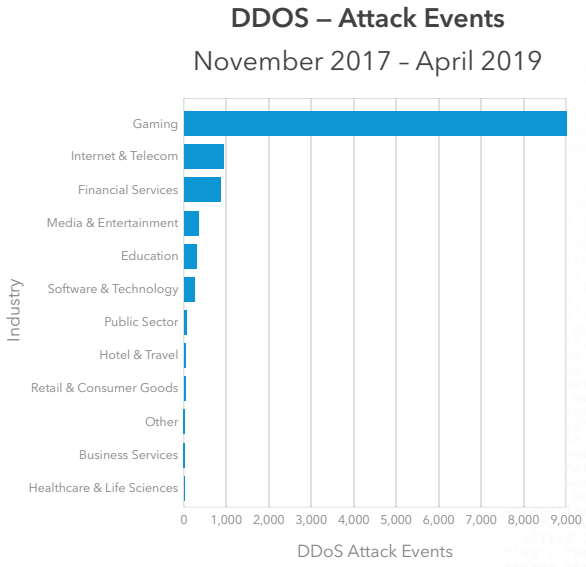

Fig. 10 – Gaming is the top DDoS target observed by Akamai with just under 9,000 attacks

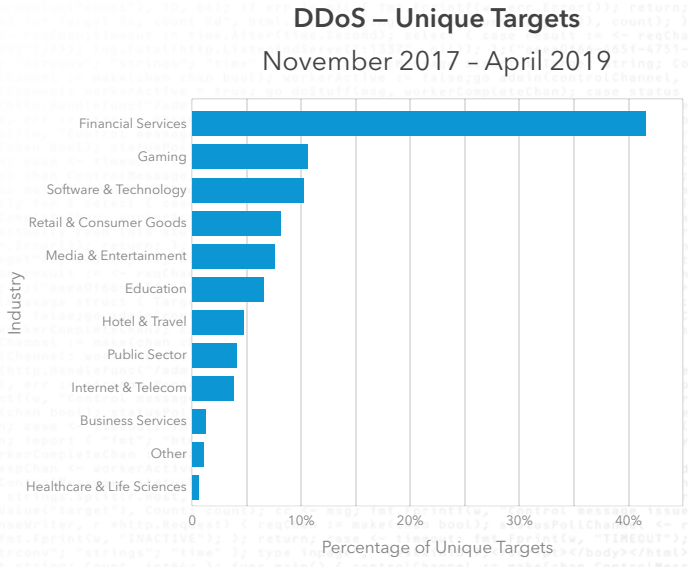

Fig. 11 – When it comes to unique DDoS targets, financial services moves to the top of the list, while gaming moves down to third

In Figure 12, we can see that the median bits per second (bps) for observed DDoS attacks against financial services is higher when compared with other industries. So not only does the financial services industry have the most unique targets, it also seems to be taking on the most malicious traffic. When compared with other industries, the median packets per second (pps) is significantly higher in the financial services sector (Figure 13).

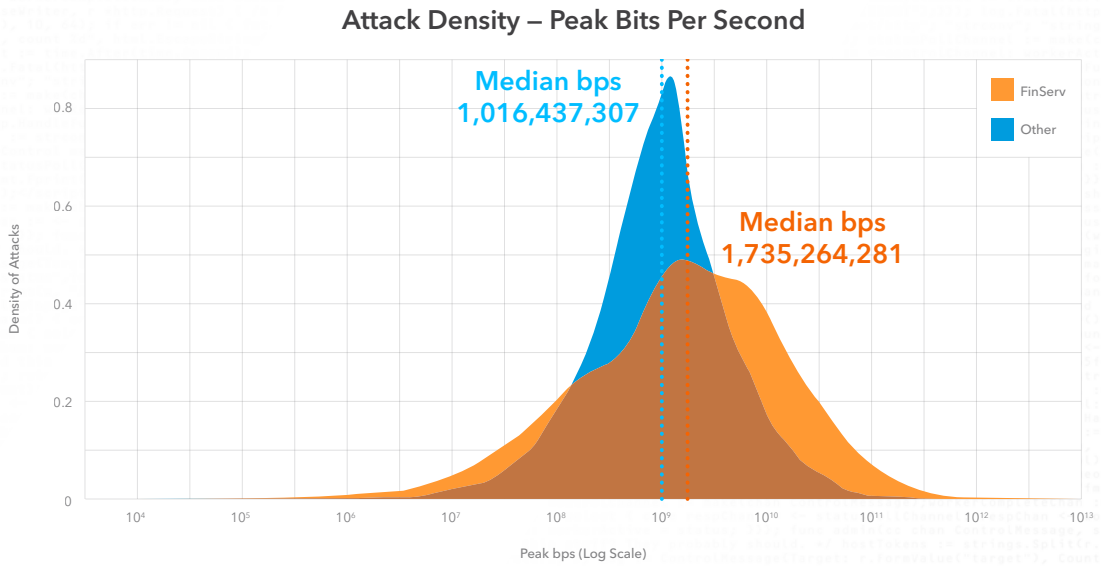

Fig. 12 – Tracking bps, financial services seems to be getting hit harder than other industries *(Please note, the area under the curve sums to 1, and the area under an interval gives the proportion of attacks that fall within that interval)*

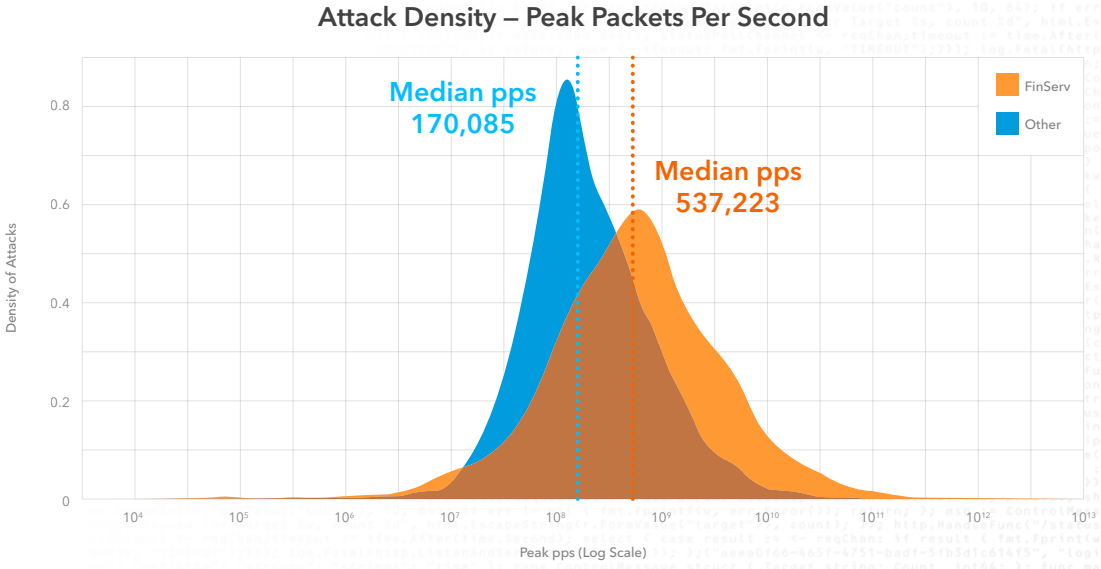

Fig. 13 – The observed median pps in financial services is significantly higher than other industries

SYN Floods, RESET Floods, TFTP Floods, and TCP Fragment Floods are the most commonly launched DDoS attack types against the financial services space. However, the industry does see a number of attack variations, meaning that banks and brokerage firms need to stay on their toes and focus on a range of defenses, instead of just one or two common attack types.

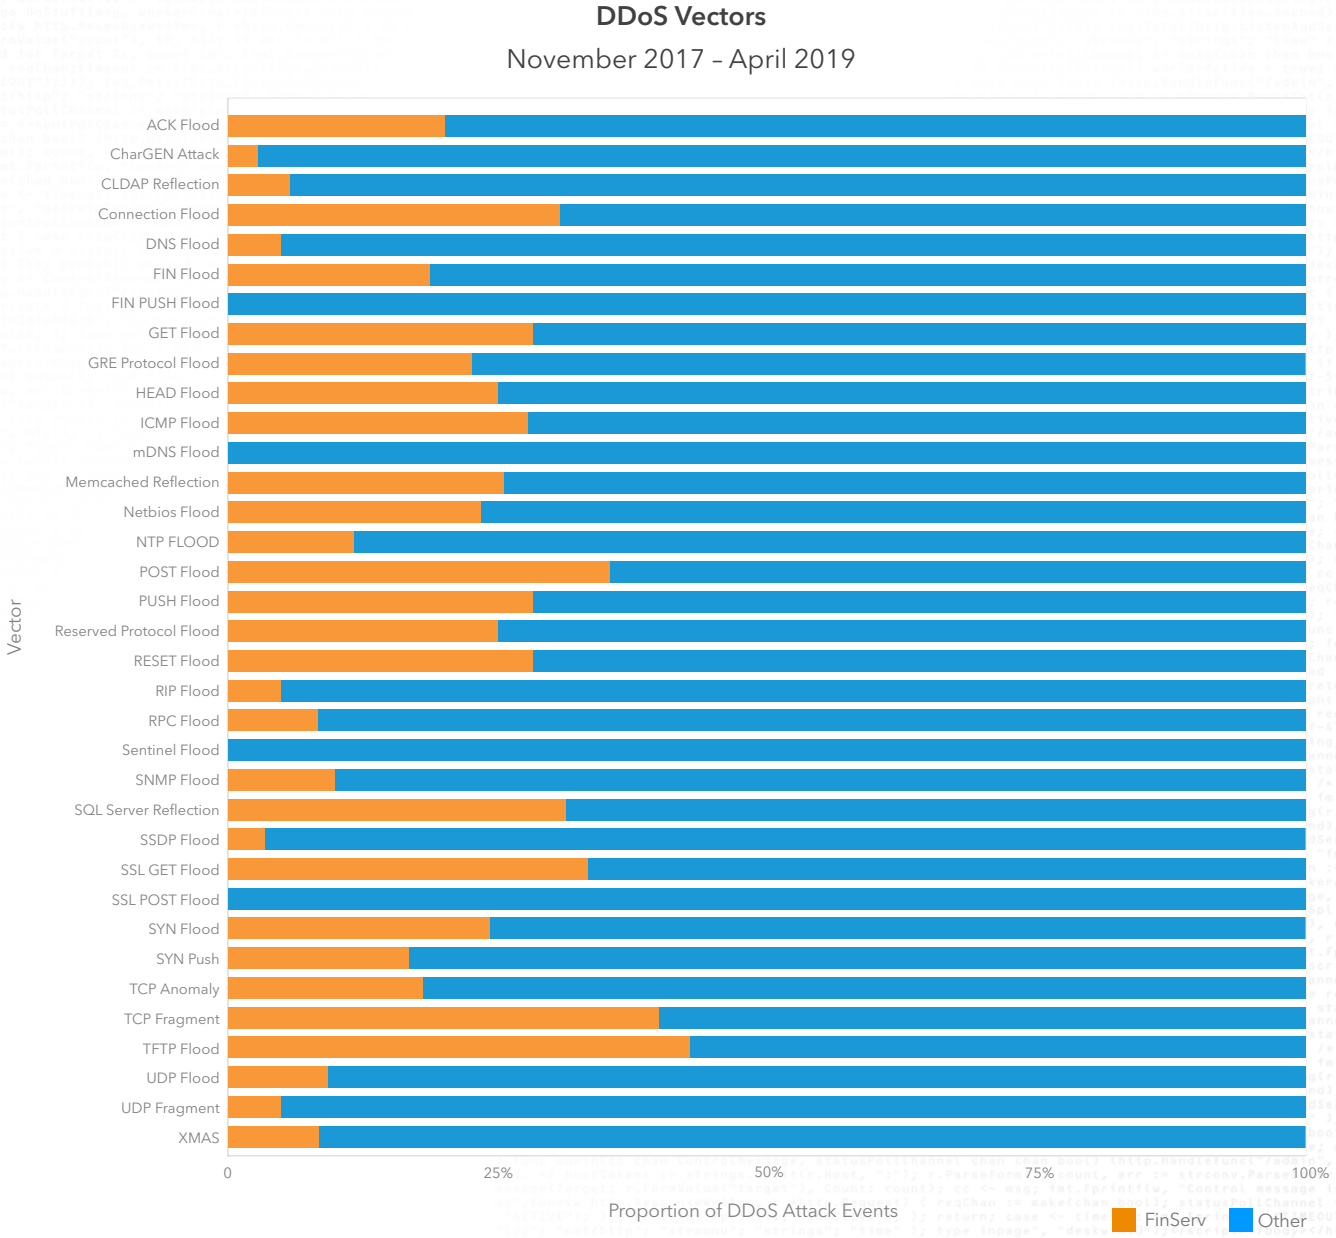

Fig. 14 – A breakdown of DDoS attack types sorted by financial services vs. all others

APIs and OFX

Criminals target authentication mechanisms; for the financial services industry, that is typically an API process or login application. Financial institutions use the Open Financial Exchange (OFX) protocol to handle data, either among themselves or to deliver the data to a third-party application.

OFX has been around since 1997. While the standard has moved up to version 2.2, many organizations are still processing data on the older, less secure, 1.x version. In 2006, OFX version 1.0.3 added basic multi-factor authentication, such as follow-up questions like a customer's mother's maiden name, birthplace, first job, etc. While it was a step in the right direction, it wasn't a foolproof protection scheme.

Direct Connect access (e.g., client-to-bank) via OFX only requires a username and password in most cases, with transactions conducted through an API. A criminal pulling a customer's details via Direct Connect would get everything needed to conduct financial fraud, including monetary theft or identity theft.

Criminals understand this, and take measures toward making their traffic look like a normal customer. Banks are starting to get savvy when it comes to picking the anomaly out of the masses. However, the slow adoption of OFX 2.2 is concerning, considering the added security it gives to financial institutions, particularly where authentication data aggregation is concerned, thanks to the adoption of Open Authorization (OAuth) and tokenization.

Using 14 days of data, Akamai examined the usage of OFX based on the traffic we have visibility into, and the results were interesting.

We observed 21,962,978 login attempts; of those, 33% (7,379,074) represented failed logins. Granted, this is a small sample set and includes just seven unique customers. It is also worth noting that the data itself doesn't label a login as malicious, only as "failed" or "successful." However, when examining the failed logins, there are some clear patterns that stand out and will raise a few eyebrows.

Overall

About 90%, or 19,706,106, of the logins (failed and successful) observed over the 14 days were over OFX v1.x, and the remaining 2,256,872 were over OFX 2.x. Where things get interesting is that 37% of logins via 1.x were failures, compared with 2.x, which hovered around 0.2%.

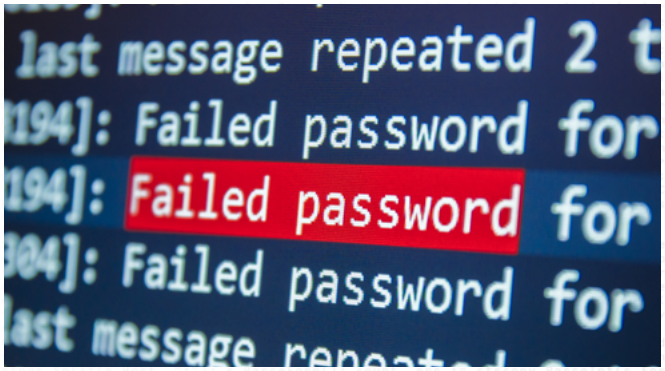

Akamai observed over 21 million OFX login attempts. Of those, 33% represented failed logins.

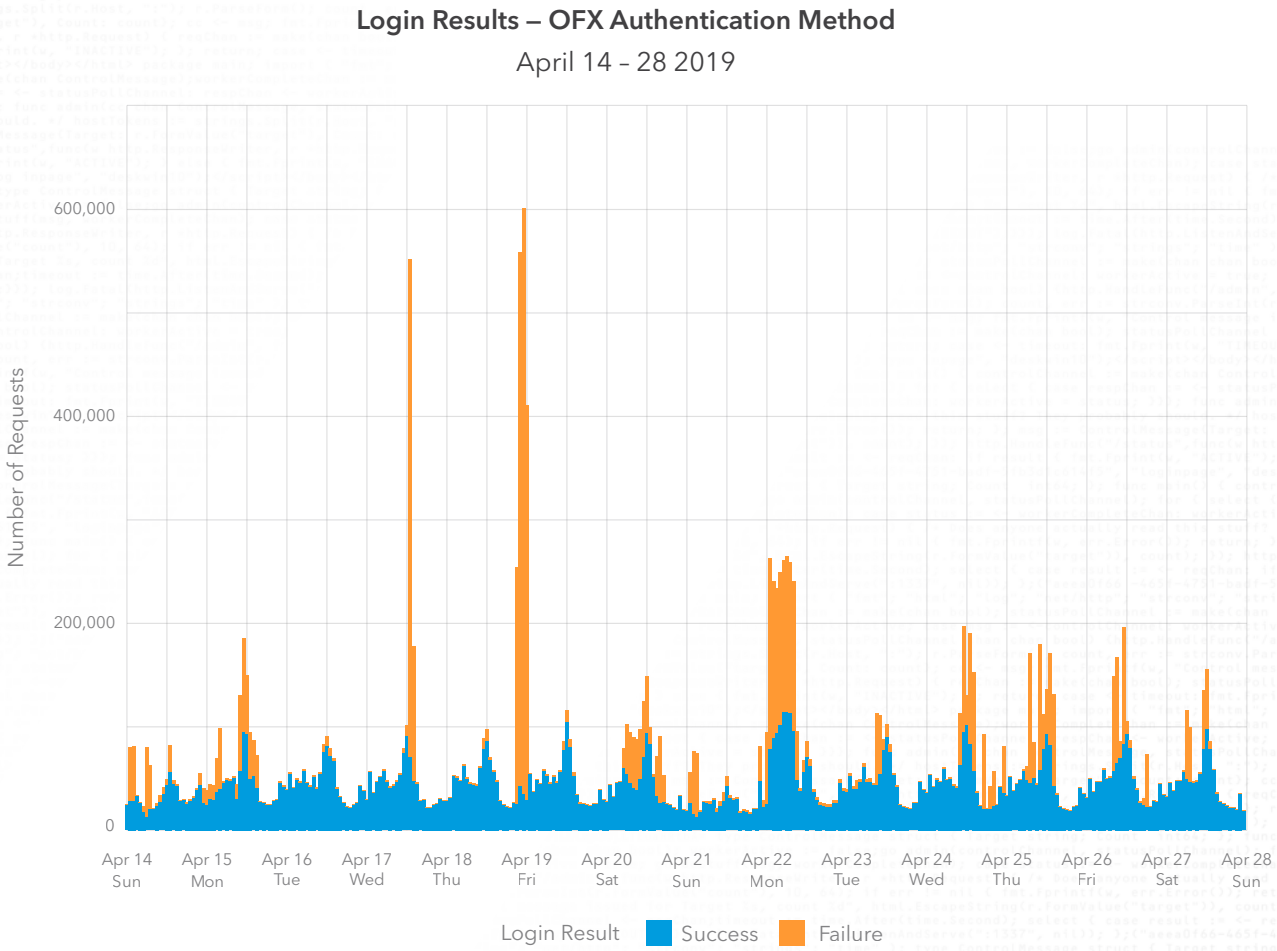

Fig. 15 – A brief look at daily OFX traffic with both failed and successful logins

The image in Figure 15 shows what we expect successful logins to look like, with regular spikes during the workday and fewer requests on the weekend. There's a small but consistent number of failed logins each hour. The median hourly failure rate is just under 3%. In the last few hours of April 18, the failure rate peaks at 94%. Based on our experience monitoring billions of credential stuffing attempts over the last year, we suspect these failed logins were due to credential stuffing attacks, but other insights into the data support this educated guess.

Failures and Successes

There were 1,896,530 unique usernames observed during this 14-day window, and 29.9% of those showed a failure rate of 0%. The mean is 17 successful logins per user, the median is 12, and the max was 5,225. It is unclear why the max was so high.

Based on the information available and considering the customers, we are making an educated guess that the max was this high because of automation and third-party connections, which are considered normal in the financial services industry.

When we examine the failed logins, 65.74% is the proportion of usernames with a failure rate between 90%–100%. We saw a mean of 4 failed logins per user, a median of 1, and a max of 4,601. Again, it isn't clear why the max is so high. However, outside of old credentials, when compared with other data points and the fact that Akamai inspected a number of OFX attacks in April 2019, this number matches credential stuffing at scale.

Credential stuffing at scale is when a service is targeted by an AIO application, and the attacker loads up several proxies to make several simultaneous connections (called threading), and their lists of usernames and passwords are tested at a steady pace so as not to draw attention to themselves.

Since the data is showing us unique usernames, and we know some credential stuffing lists will take a username and attempt multiple passwords or password variations during targeted attacks, the large number of failed logins starts to paint a bleak picture.

To visualize this, Figure 16 offers a breakdown. The first bar represents all users with a 0%–10% failure rate. The last bar is all users with a 90%–100% failure rate.

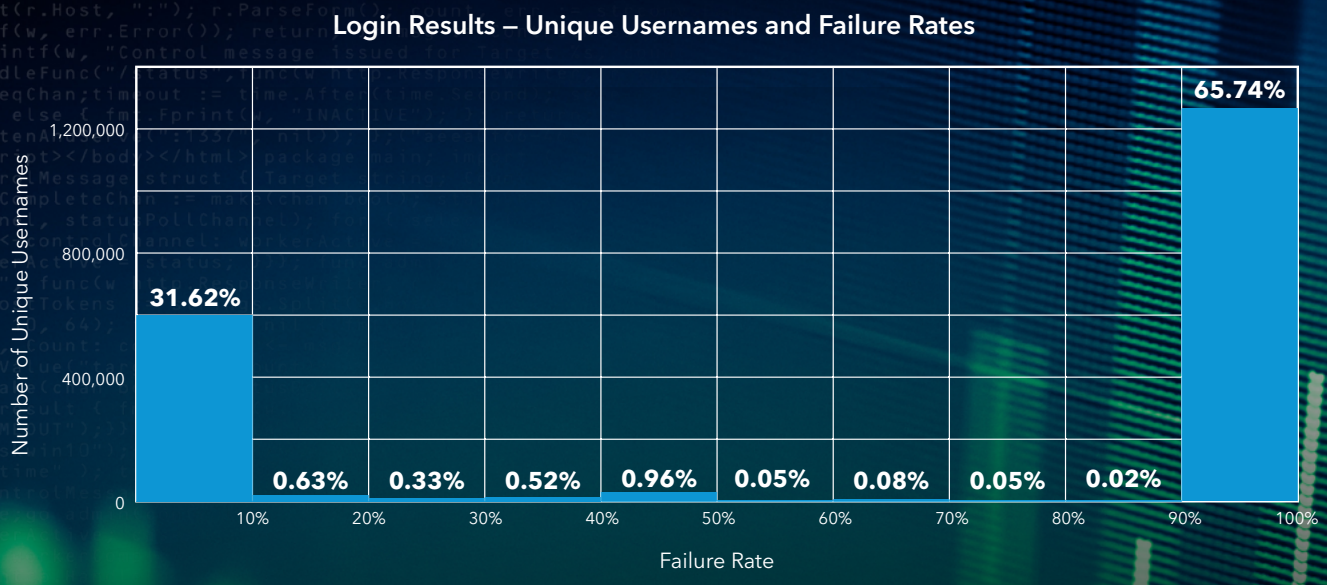

Fig. 16 – A breakdown of login failures and successes by username

Looking at other login data and measuring by success or failure, in Figure 17, we can see 1.5 times more unique IPs, when comparing failed logins with successful logins; 22 times more unique Transport Layer Security (TLS) fingerprints, and twice the number of usernames.

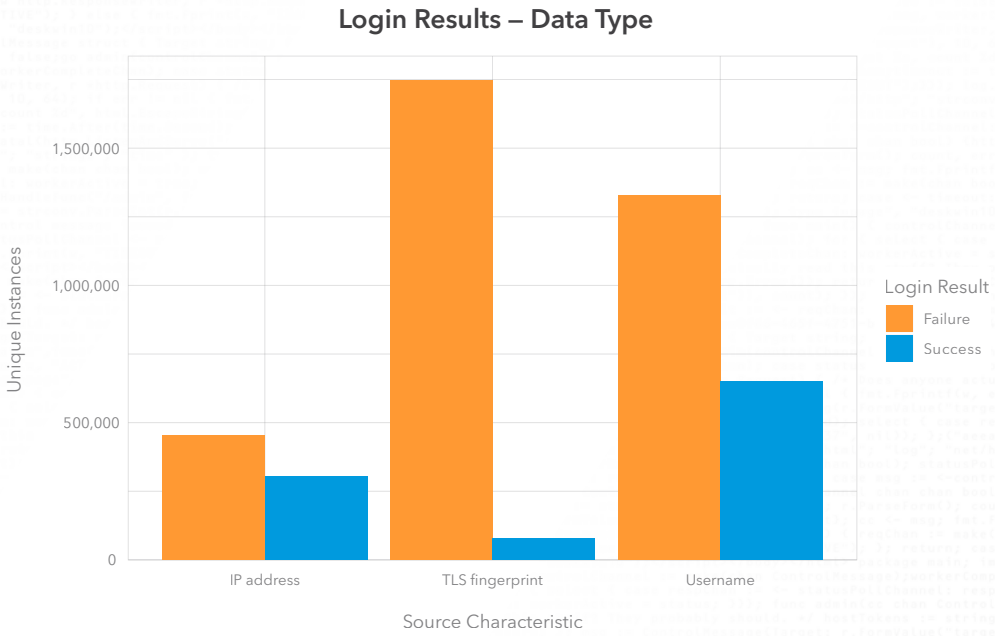

Fig. 17 – A breakdown of login data measured by fail or success, and data type

The unique TLS fingerprints were another indicator of a credential stuffing attack. More than 95% of all TLS fingerprints exhibited a failure rate of more than 90%; 1,672,519 fingerprints had a failure rate of 100%.

Earlier this year, Akamai published research on Cipher Stunting, where criminals randomize TLS fingerprints in order to avoid detection. We believe this data is an indicator of such actions, and proof that the financial services industry is the key target for these types of attacks.

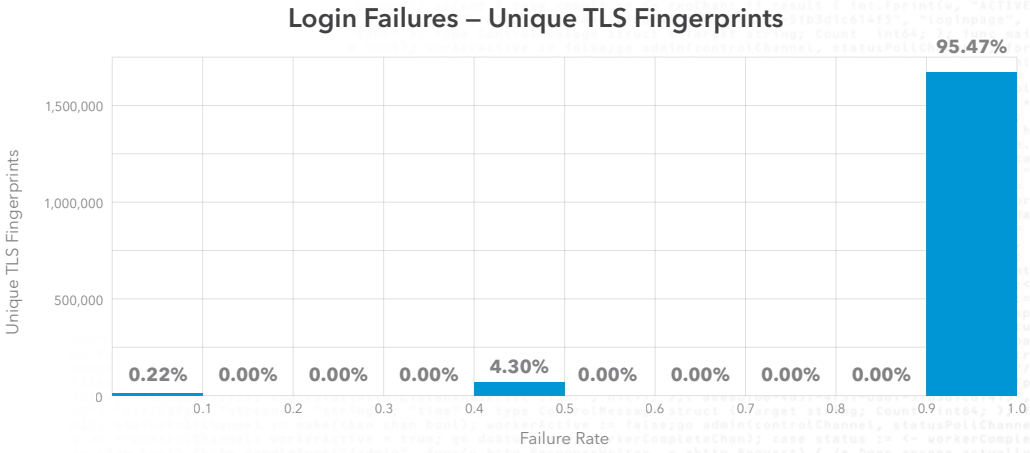

Fig. 18 – A breakdown of TLS fingerprints

Criminal Economy

Once the attacks end, assuming they are successful, the amount of data available to criminals will vary depending on the organization they targeted. The financial services industry as a whole is a treasure trove of information and resources for a criminal — and everything has a value associated with it.

A large part of the criminal economy is fueled by financial services data. Prices rise and fall, but the steady stream of financial information available in these markets is almost endless. As this report was being written, one darknet forum was flush with new accounts that could be leveraged as bank drops in the United States.

Bank drops are packages of data and services that can be used to open accounts at a given financial institution. Some sellers will create and develop the bank drop for the buyer and provide required details or services, while others will just provide the information needed, and the buyer is on their own.

Each bank drop will include a person's stolen identity (sometimes called "fullz"), including full name, address, date of birth, Social Security number, driver's license data, credit score details, and access to a secure Remote Desktop Protocol (RDP) connection for one month. The RDP is clean, meaning it has never been used for anything malicious or scamming other than the development of the drop, and the location of the computer will match the stolen identity and the bank branch.

Currently, drops at two major banks were selling for $150, $200, and $250 per account. The price variation is due to the extra services offered. The lowest-priced package consists of just the newly created drop, while the other two include an additional mail drop where account-related communications and bank cards can be delivered.

The packages also include account username and password, security questions and answers, routing and account numbers, and either a local VoIP or SIM card number that can be used to receive an SMS or to call the bank for verbal confirmations if needed.

Another seller had a cache of drops available for one of seven different consumer banks (Figure 19) at prices ranging from $300 to $400. The lowest price point is for a random bank drop, while the others offered the buyer their choice of bank. The top tier includes choice of bank, as well as an active and verified web-based payment processing account.

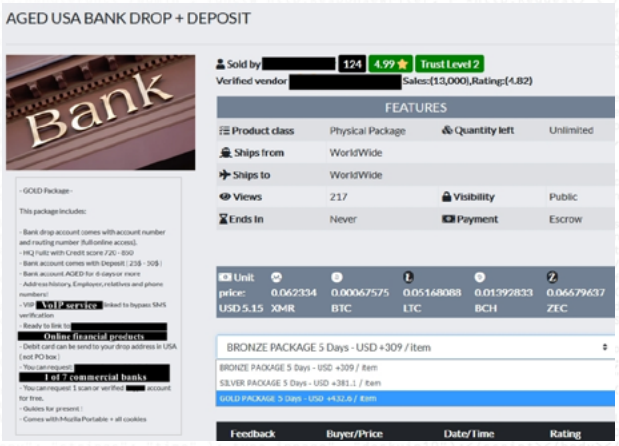

Fig. 19 – A darknet ad selling commercial bank drops

The package includes a detailed set of identity information, as well as a starting deposit of $25–$50. However, the hard selling point for this offer was that the accounts were pre-aged for a number of days and created for the buyer by the seller.

In addition, the second seller's identity kit also includes address history information, previous employment records, information on the person's relatives, and all of their contact details. For technical and security concerns, the offer also included a copy of the portable browser and all relevant cookies that were assigned during account creation.

In the event the drops were detected by the bank and closed, assuming voice verification didn't work to restore access, both sellers were willing to offer replacements free of charge under certain circumstances.

The aim of these drop accounts is to provide a place to store funds that are cashed out. Cashing out is a term used to explain the process of emptying compromised bank accounts, maxing out stolen credit cards, or the process of shifting digital currency into cash. Either way, the criminals need a bank to safely store and use their funds. But over the years, this has become increasingly more difficult.

Cashing out is a labor-intensive process for criminals, and comes with a high degree of risk. In fact, the risk is so high that many criminals use cutouts or money mules to do most of the dirty work. They advertise these positions in a number of places, usually marketing them as work-from-home opportunities.

A criminal with a series of drops can use them to launder money. However, because of strict reporting and security rules at many financial institutions, this aspect of a drop isn't as common as it once was. Instead, criminals will funnel their money into digital currency and use web-based cleaning services before slowly cashing out into a drop account.

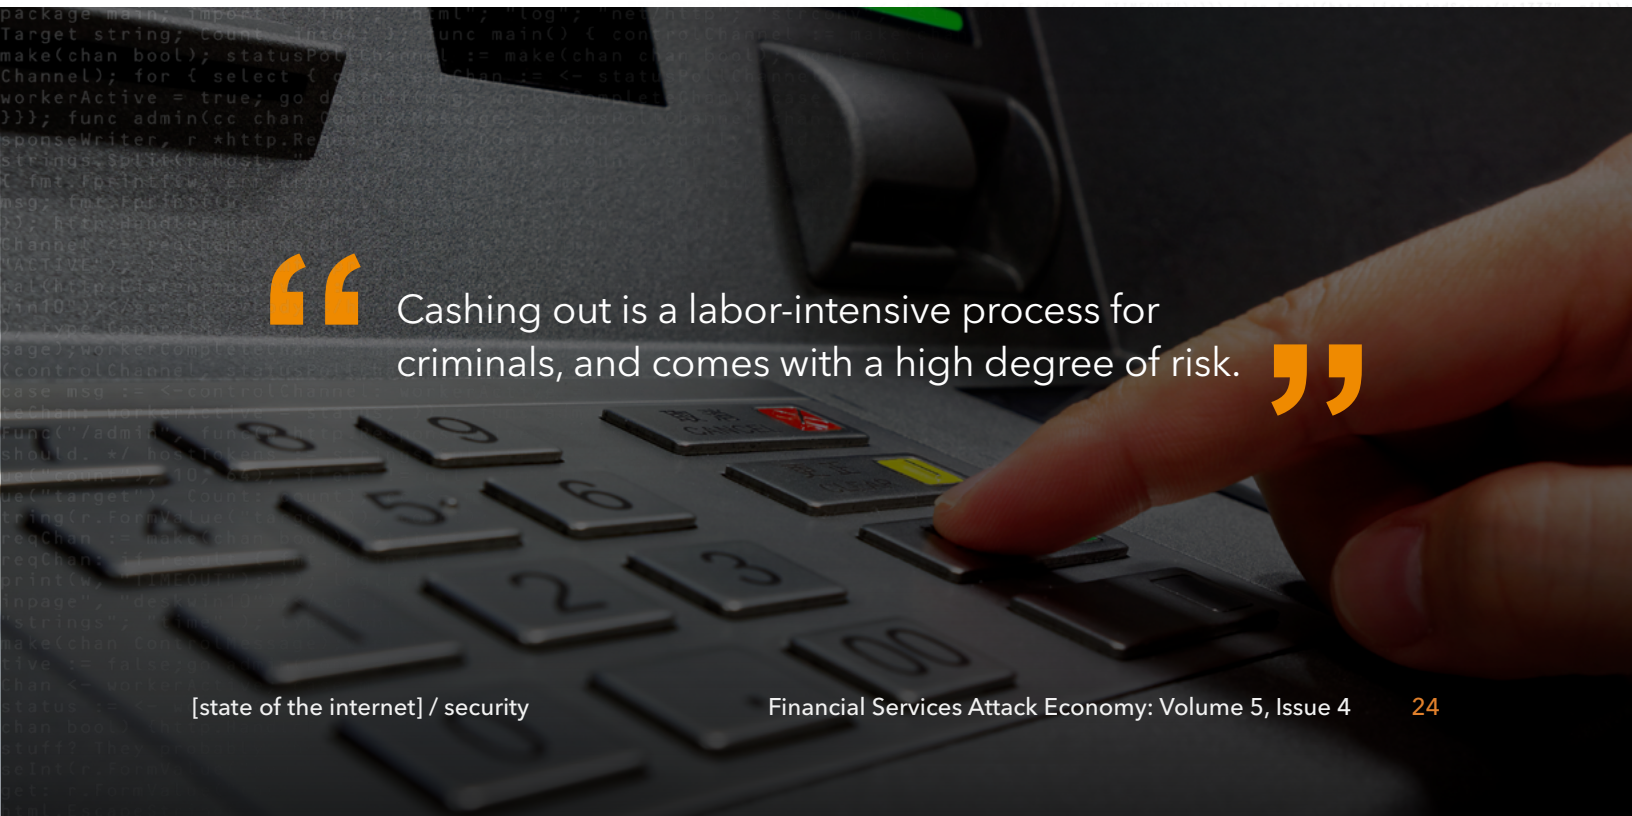

Financial Services Attack: A Personal Story

Steve Ragan

Sr. Technical Writer

Editor, State of the Internet / Security

In the financial services industry, credential stuffing is a real threat. It impacts the organizations dealing with financial matters, and it affects their customers. When credential stuffing attacks are successful, they can sometimes generate sensational and embarrassing headlines, but the story that is rarely told – at least in the public eye – is what happened to me.

I was a credential stuffing attack target, but fortunately for me, I wasn't a victim. The reason I wasn't a credential stuffing victim can be explained in two parts.

The first, and most important reason, is because my bank stopped the attack. On April 27, 2019, at about 4:30 a.m., I received an email from my bank telling me that my services online and via phone were suspended "... due to multiple attempts using [incorrect] credentials ..."

I knew something was up. Less than 12 hours earlier, one of my social media accounts triggered with password reset requests. Around the same time, several other accounts also warned me of failed logins or attempted to get me to enter verification codes. I'm not a fan of coincidence, and this was too many attempts, spread across too many accounts, to be a basic brute force attack. Someone was running a list.

After verifying ownership of my banking account and changing my password (just to be safe), my access was restored. This lasted for just over 24 hours, after which the account was disabled again for the same reasons. A second round of verification and password resets, and things were good to go.

I got lucky. My bank pays close attention to failed logins and doesn't hesitate to disable accounts. If the attack had succeeded, I would have lost an account that contained money used to pay a car loan, as well as any personal information associated with it.

Speaking to others who use the same bank, the entire week was a rocky one. Each person I spoke to said they also had their accounts disabled for the same reasons.

Aside from each of us sharing the same bank and placement on credential stuffing lists being circulated by criminals, each of us avoided being victimized by the attacks because we used unique passwords and took advantage of all of the bank's account protections and notifications.

The second reason I wasn't a credential stuffing victim? I knew the password that the criminals attempted to use, and where it came from. The password was a tagged one, chosen to be weak and tied to the domain where it was used (e.g., domain123). At the time, I just didn't trust the website where I registered the password, and I didn't care if the account was compromised. Earlier this year, when a flood of username and password combinations appeared on the web, I discovered my tagged credentials in the dump.

I expected this sort of attack to happen eventually, and I was confident that it wouldn't succeed. Hubris? Maybe, but I don't mix my passwords. My banking password isn't related to my email passwords, and they're not related to social media, etc. Honestly, I don't even know the majority of my passwords. I use a password vault to manage all of the important ones.

But still, I freely admit I got lucky. The problem is, luck runs out – even for security professionals. So I remain vigilant, watch my accounts, and try to stay ahead of the game when it comes to account security. That's all any of us can do, and one can only hope that's enough.

26 Looking Forward

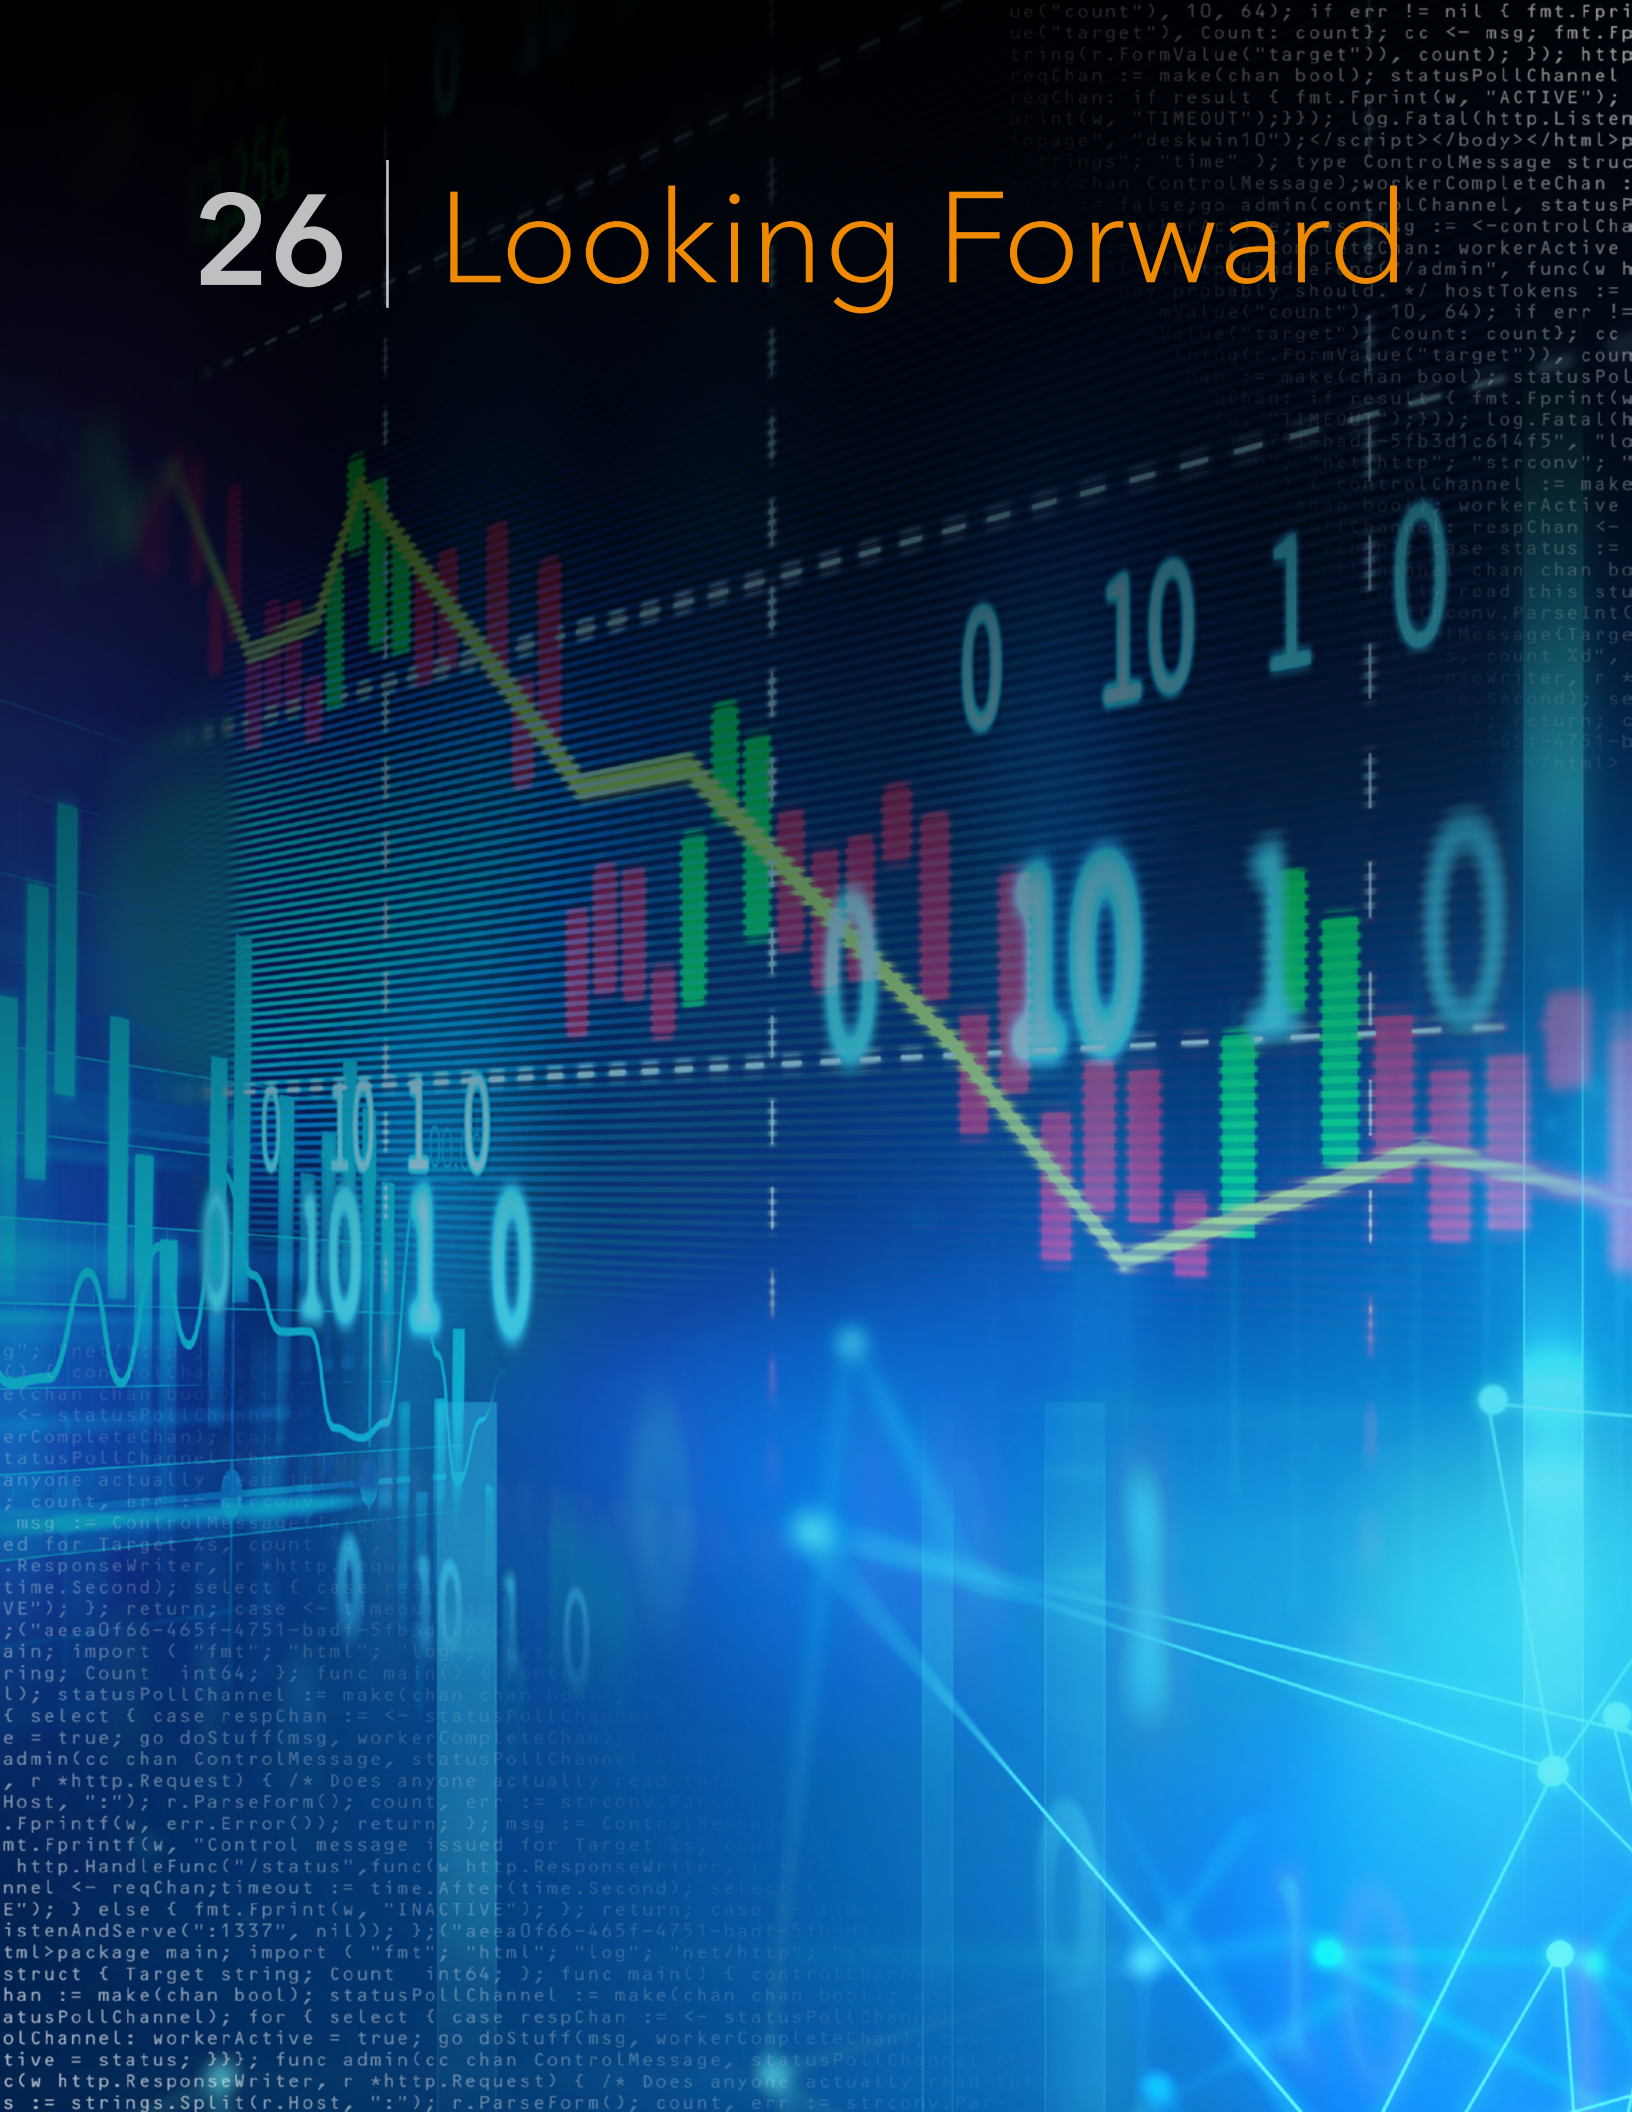

The circular battle between banks and criminals isn't going away, since it is tied to the core of the criminal economy. The banks that criminals target are the very same ones they'll use to handle their ill-gotten profits.

The financial services industry spends billions on security each year; in doing so, they've made the criminal economy come out from the shadows. In the past, fraudulent accounts would exist for weeks before they were detected. Now, those same accounts can disappear within hours, which is part of the reason why aged drop accounts sell for such high amounts. The risks and pressures associated with these types of crimes are getting to be too much.

It takes a layered approach to attack the financial sector. From credential stuffing and phishing, to the associated goal of identity theft and DDoS attacks, there is no assured success in any of these criminal endeavors.

For example, in April of 2019, a Pennsylvania man was sentenced to 64 months in prison and three years of supervised release for money laundering and identity theft. He managed a carding and cash-out scheme for more than two years, and continued to do so after initially confronted by police. Using the same types of data that criminals target when conducting phishing or credential stuffing attacks, he leveraged stolen identity information and stolen credit cards to purchase gift cards and other items that were then later resold in order to turn their value into spendable cash. In all, the individual and his brother laundered more than $218,000 during their spree.

In 2018, a California man pleaded guilty to access device fraud (i.e., using stolen and counterfeit credit cards, debit cards, account numbers, and other financial account information), as well as identity theft. According to court documents, police found several boxes containing stolen or forged documents, bank cards, and identification. During stipulation in open court, the man admitted to obtaining his victims' information via a number of means, including online account takeovers. His crime spree caused at least $558,276 in damages.

Over the years, criminals have become increasingly savvy when it comes to how they manipulate the financial industry. They have hundreds of schemes on deck to rob and steal, but the financial services industry has developed hundreds of processes to detect and stop said schemes.

Due to a growing security ecosystem and tightened regulatory controls, financial firms are constantly expanding their security acumen; criminals now have to fight and expose themselves to a growing level of personal risk in order to keep up with financial marks that are no longer the easy targets they once appeared to be.

And anything that makes a criminal's day harder is a good thing in our book.

> "In the past, fraudulent accounts would exist for weeks before they were detected. Now, those same accounts can disappear within hours…"

28 | Appendix A: Methodologies

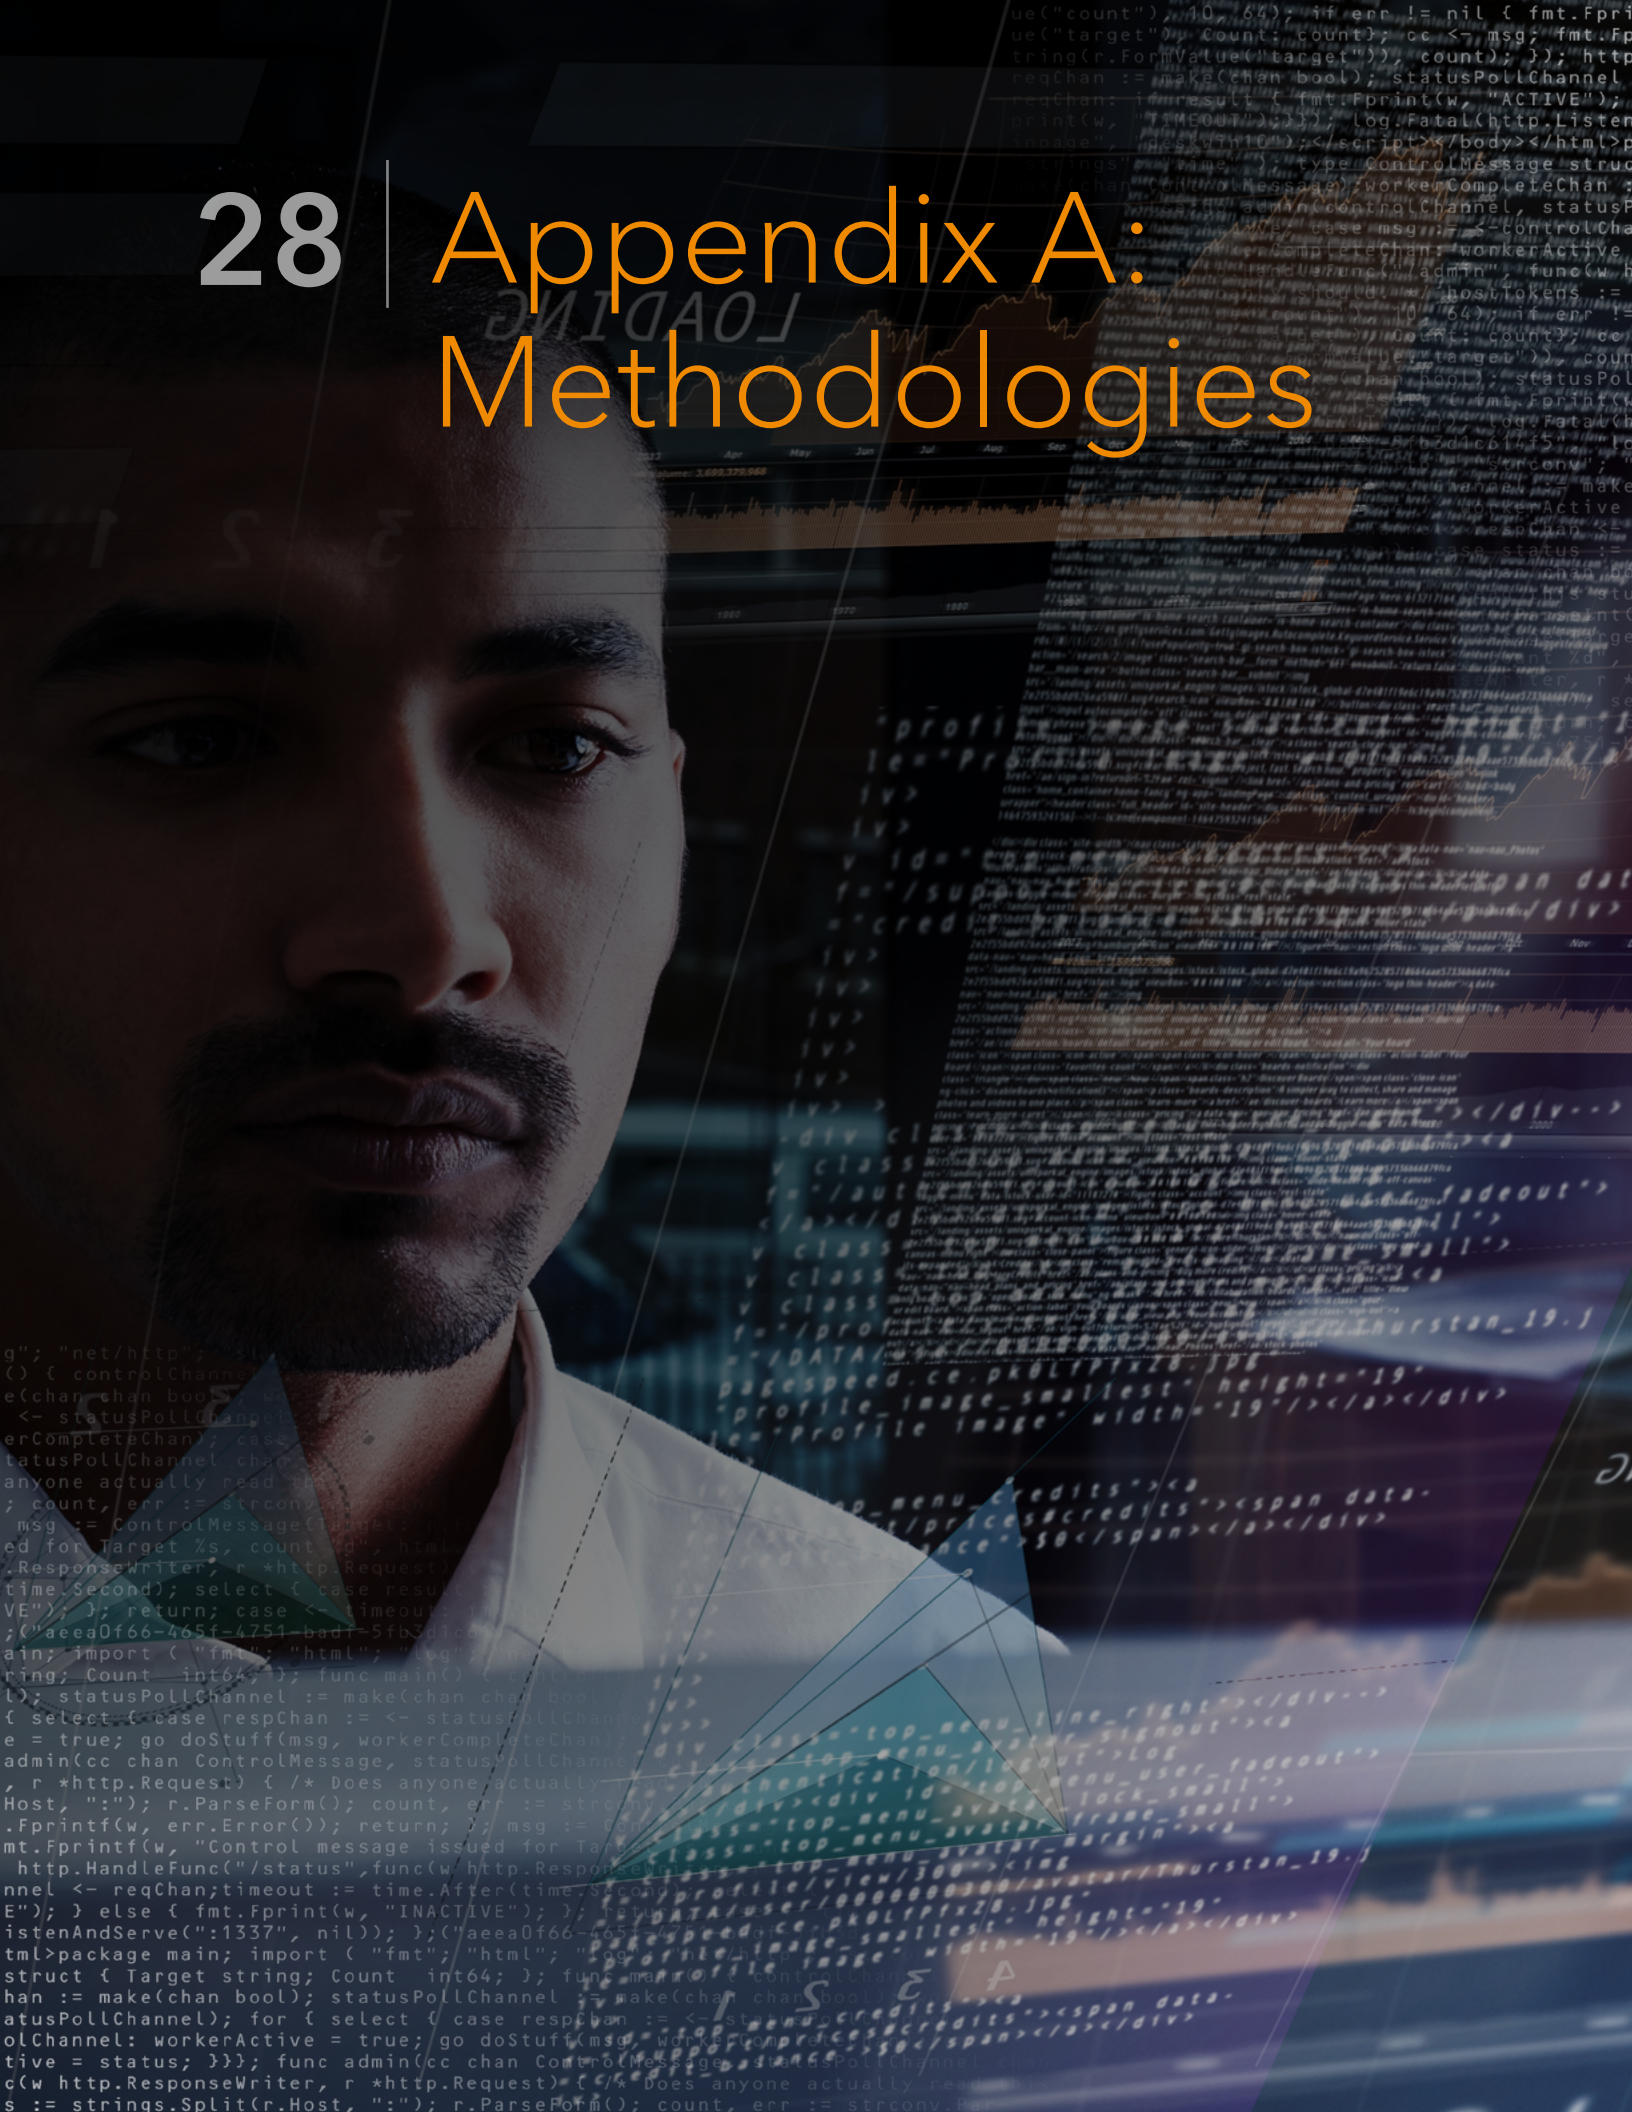

General Notes

The team that creates the State of the Internet / Security report does the best we can to make our data as clear and accurate as possible.

Web Attacks

The Akamai Intelligent Edge Platform is a network of more than 230,000 servers in thousands of networks around the world. The Kona WAF is used to protect this traffic, and the information about the attacks is fed into an internal tool called Cloud Security Intelligence (CSI). This data, measured in petabytes per month, is used to research attacks, understand trends, and feed additional intelligence into Akamai's solutions. This data represents millions of daily application layer alerts, but these alerts do not indicate a successful compromise.

The plots and tables provided in this section were limited to records between November 1, 2017, and April 30, 2019.

Credential Abuse

The data for this section was also drawn from the CSI repository. Credential abuse attempts were identified as unsuccessful login attempts for accounts using an email address as a username. In order to identify abuse attempts, as opposed to real users who can't type, two different algorithms are used. The first is a simple volumetric rule that counts the number of login errors to a specific address. This differs from what a single organization might be able to detect because Akamai is correlating data across hundreds of organizations.

The second algorithm uses data from our bot detection services to identify credential abuse from known botnets and tools. A well-configured botnet can avoid volumetric detection by distributing its traffic among many targets, using a large number of systems in its scan, or spreading out the traffic over time, just to name a few.

These records were collected between November 1, 2017, and April 30, 2019. As additional data is collected, future reporting will target a rolling two-year time frame.

30 | Appendix B: Supplemental Data

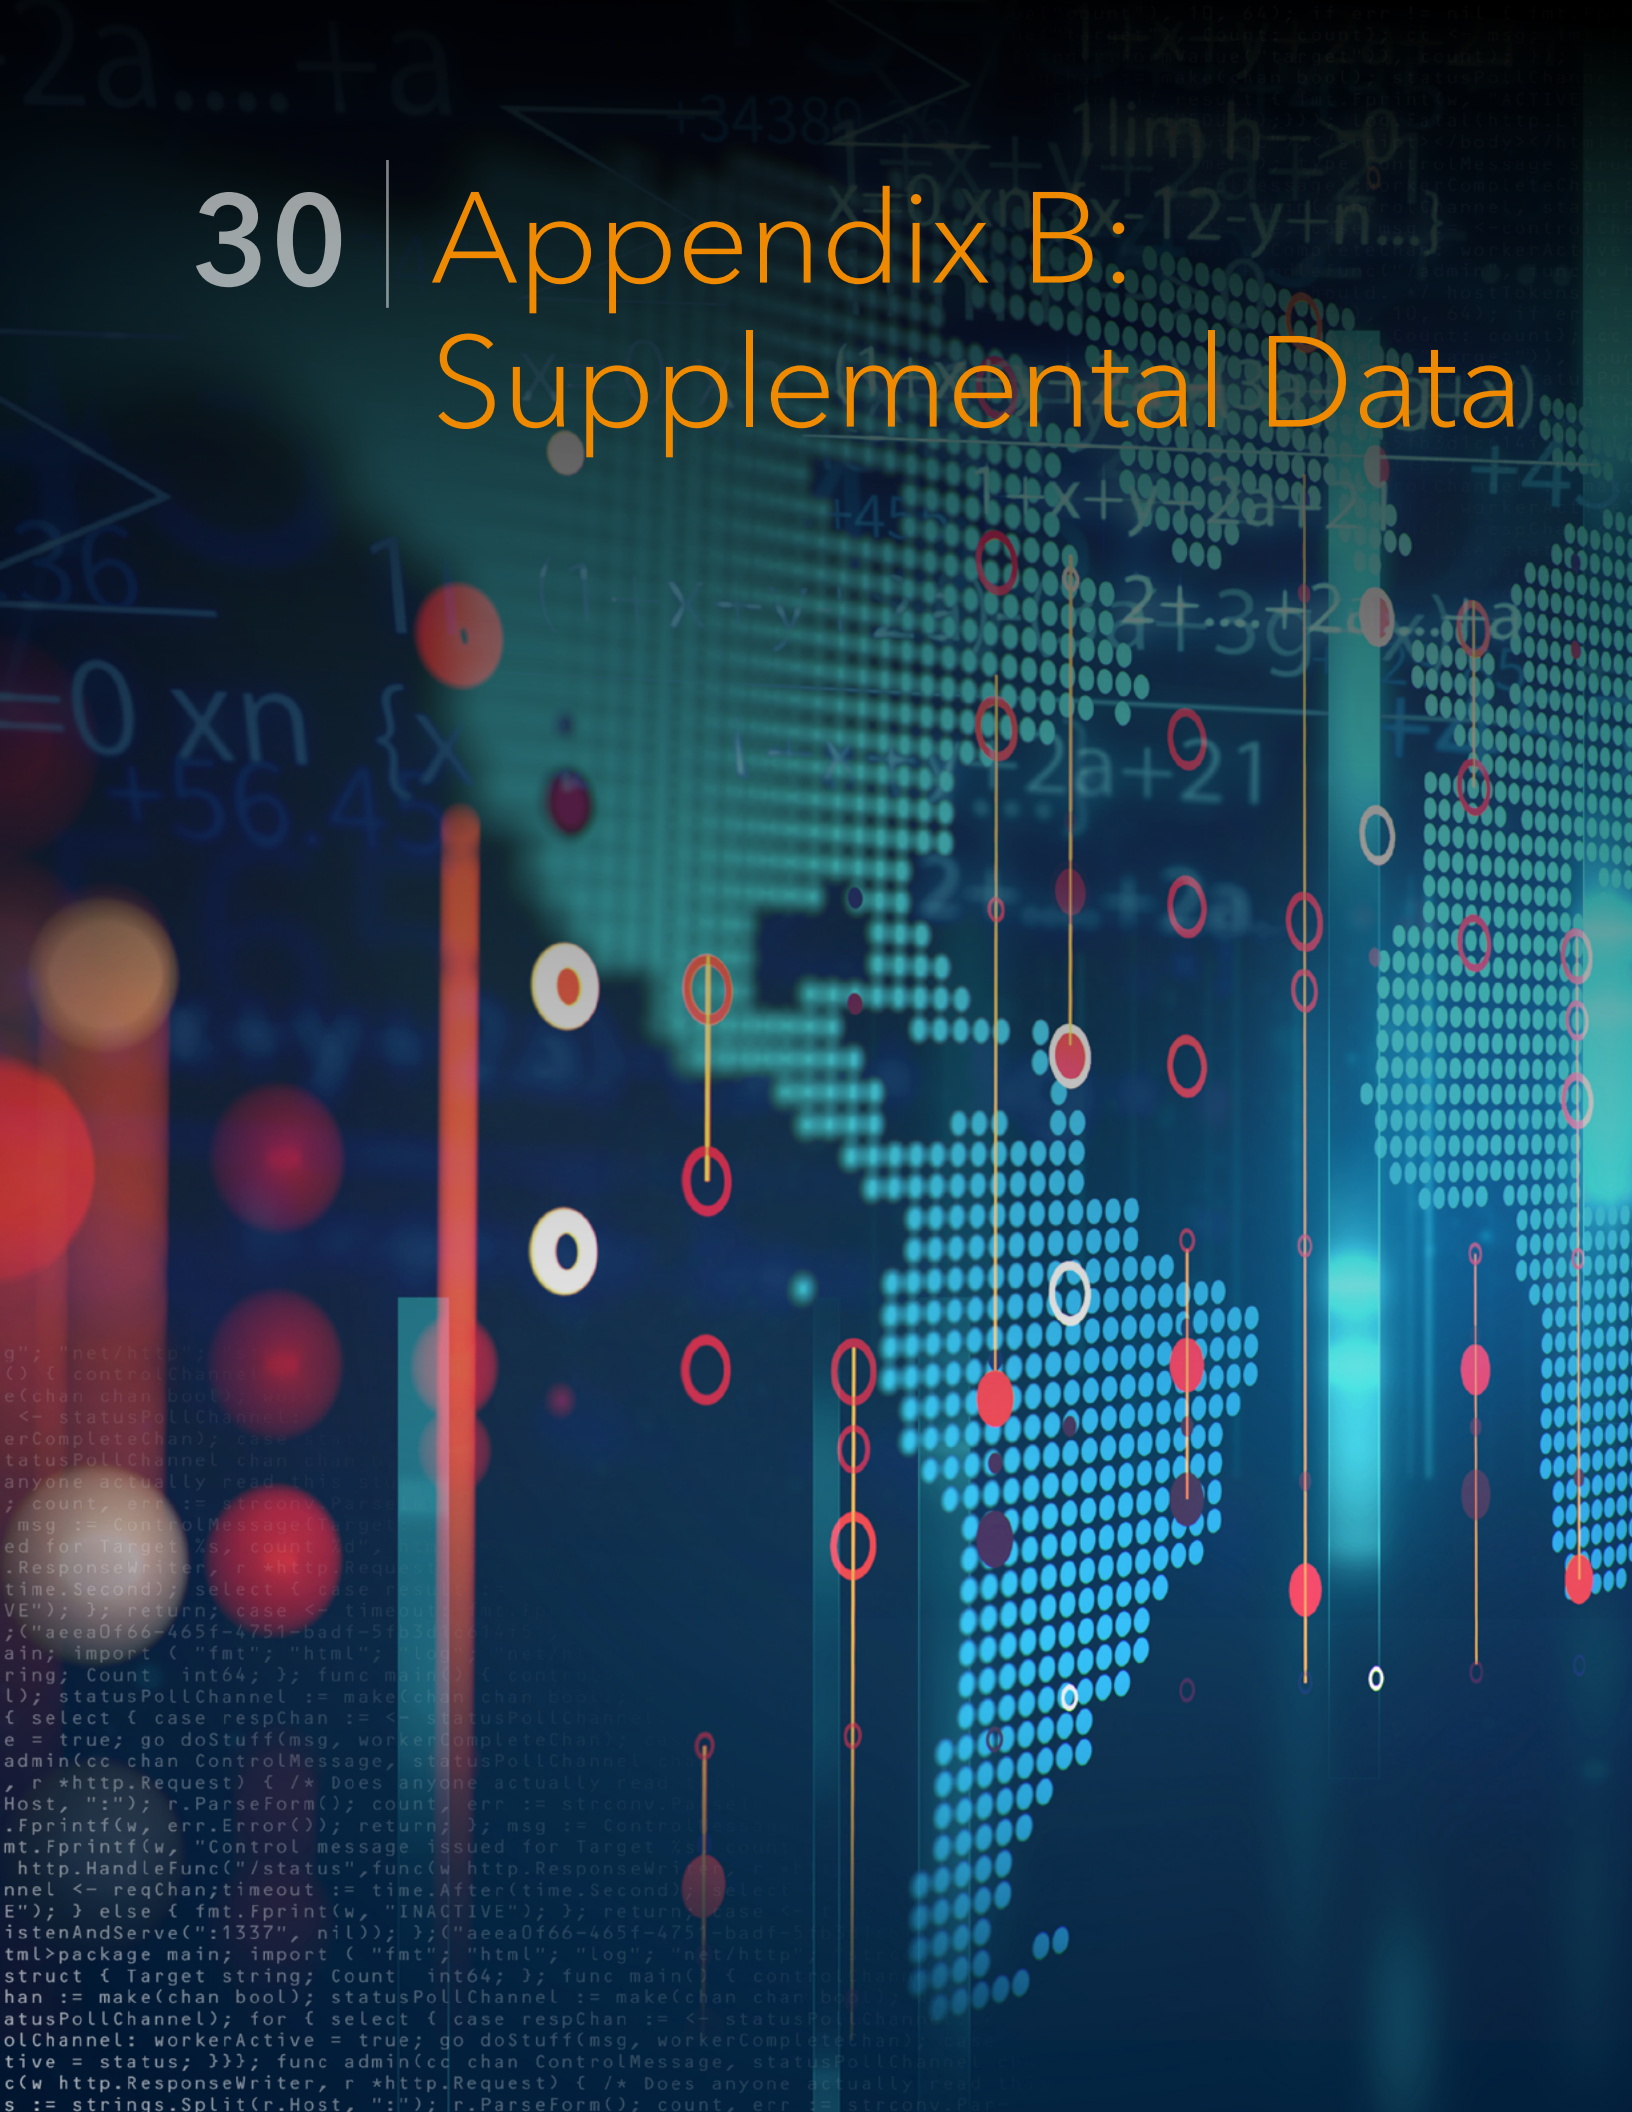

The following plots are provided to show the regional differences in some of the plots used in the main body of the State of the Internet / Security report. They are roughly divided into the Americas; the Asia Pacific (APAC) region; and the Europe, Middle East, and Africa (EMEA) region. Our goal was to show the top representatives in each region; the plots do not include the entirety of our data sets.

Credential Abuse Attack Sources – Americas

November 2017 – April 2019

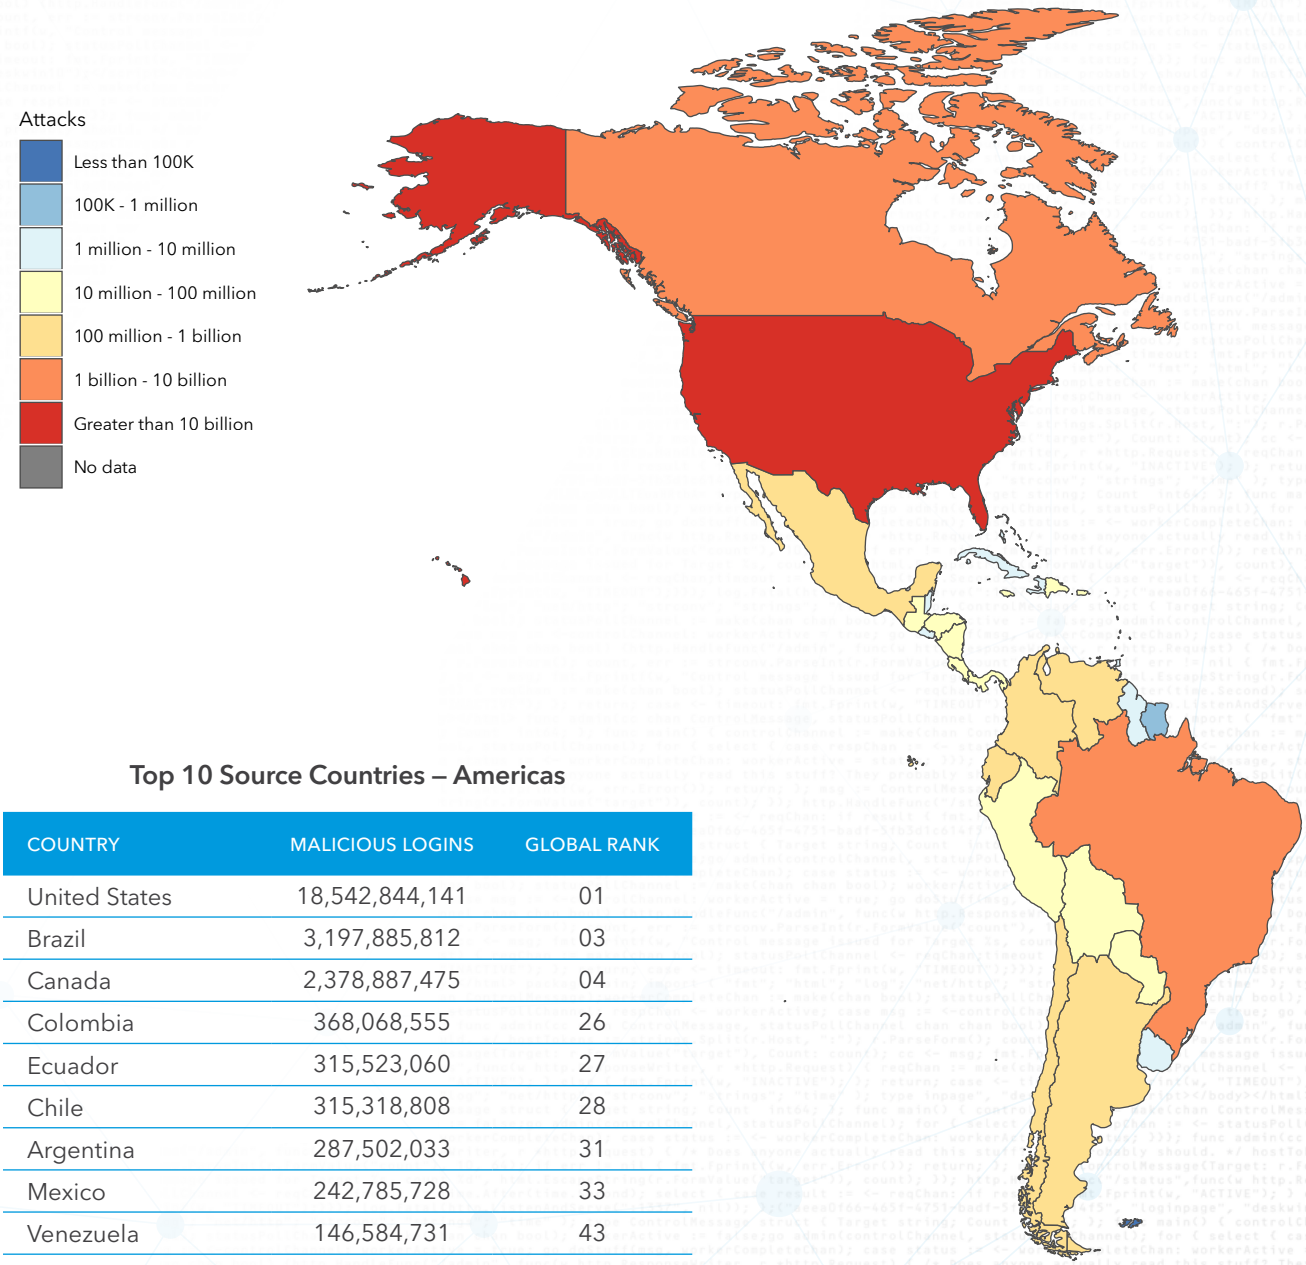

Top 10 Source Countries – Americas

COUNTRY	MALICIOUS LOGINS	GLOBAL RANK
United States	18,542,844,141	01
Brazil	3,197,885,812	03
Canada	2,378,887,475	04
Colombia	368,068,555	26
Ecuador	315,523,060	27
Chile	315,318,808	28
Argentina	287,502,033	31
Mexico	242,785,728	33
Venezuela	146,584,731	43
Dominican Republic	63,667,312	61

Credential Abuse Attack Sources – EMEA

November 2017 – April 2019

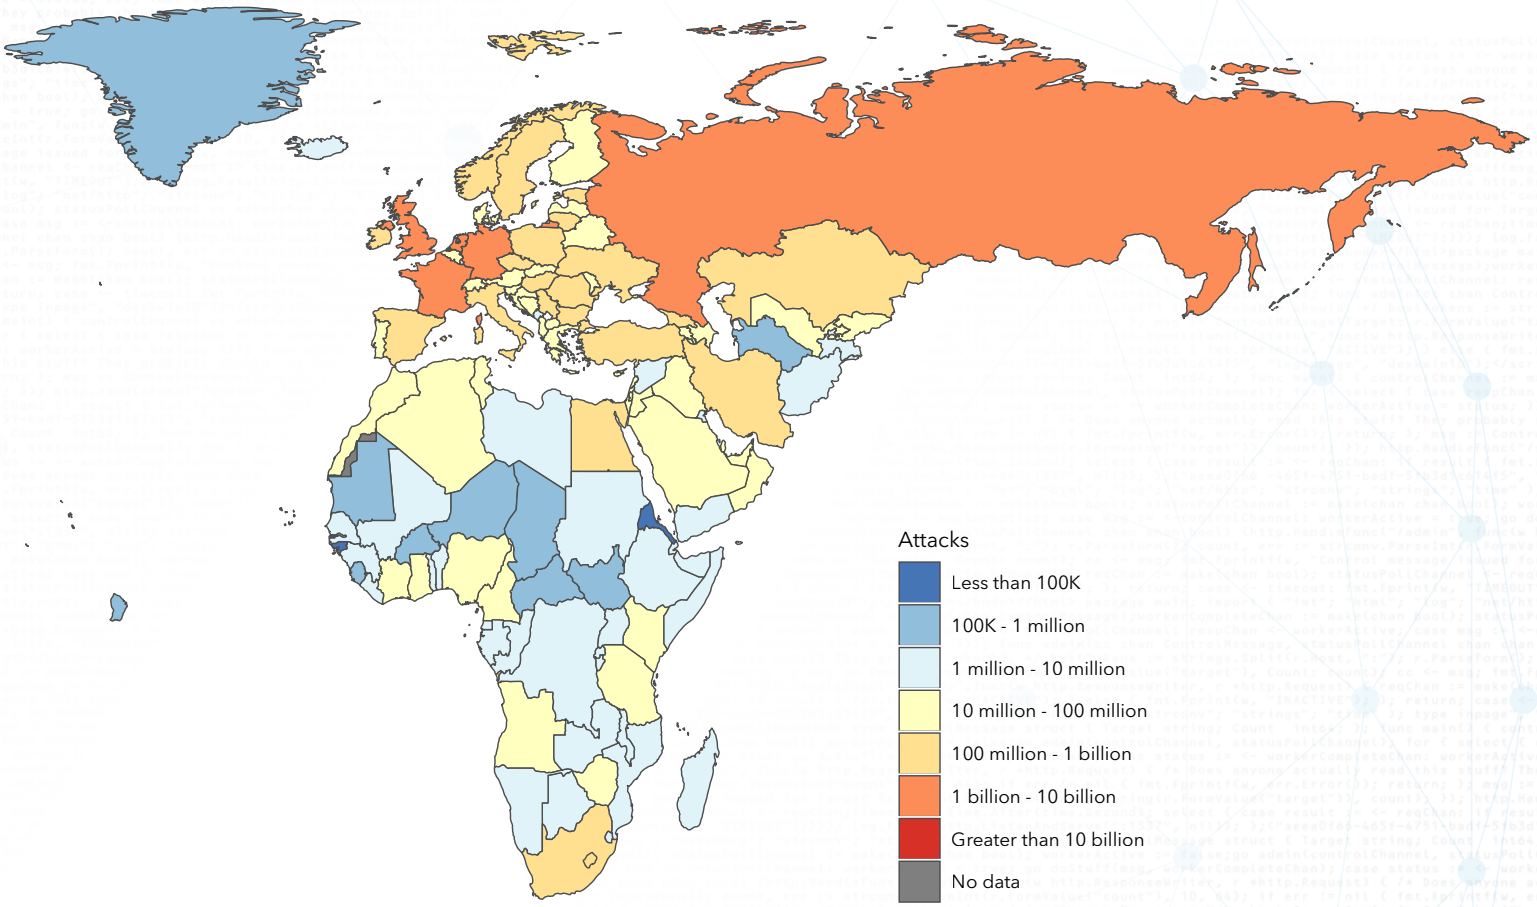

Top 10 Source Countries – EMEA

COUNTRY	MALICIOUS LOGINS	GLOBAL RANK
Russia	5,389,826,207	02
Netherlands	1,361,410,559	10
France	1,358,348,831	11
Germany	1,357,470,150	12
United Kingdom	1,086,213,915	13
Ukraine	781,992,409	16
Italy	621,338,014	17
Estonia	465,763,965	21
Poland	425,905,077	23
Spain	386,650,197	24

Credential Abuse Attack Sources - APAC

November 2017 – April 2019

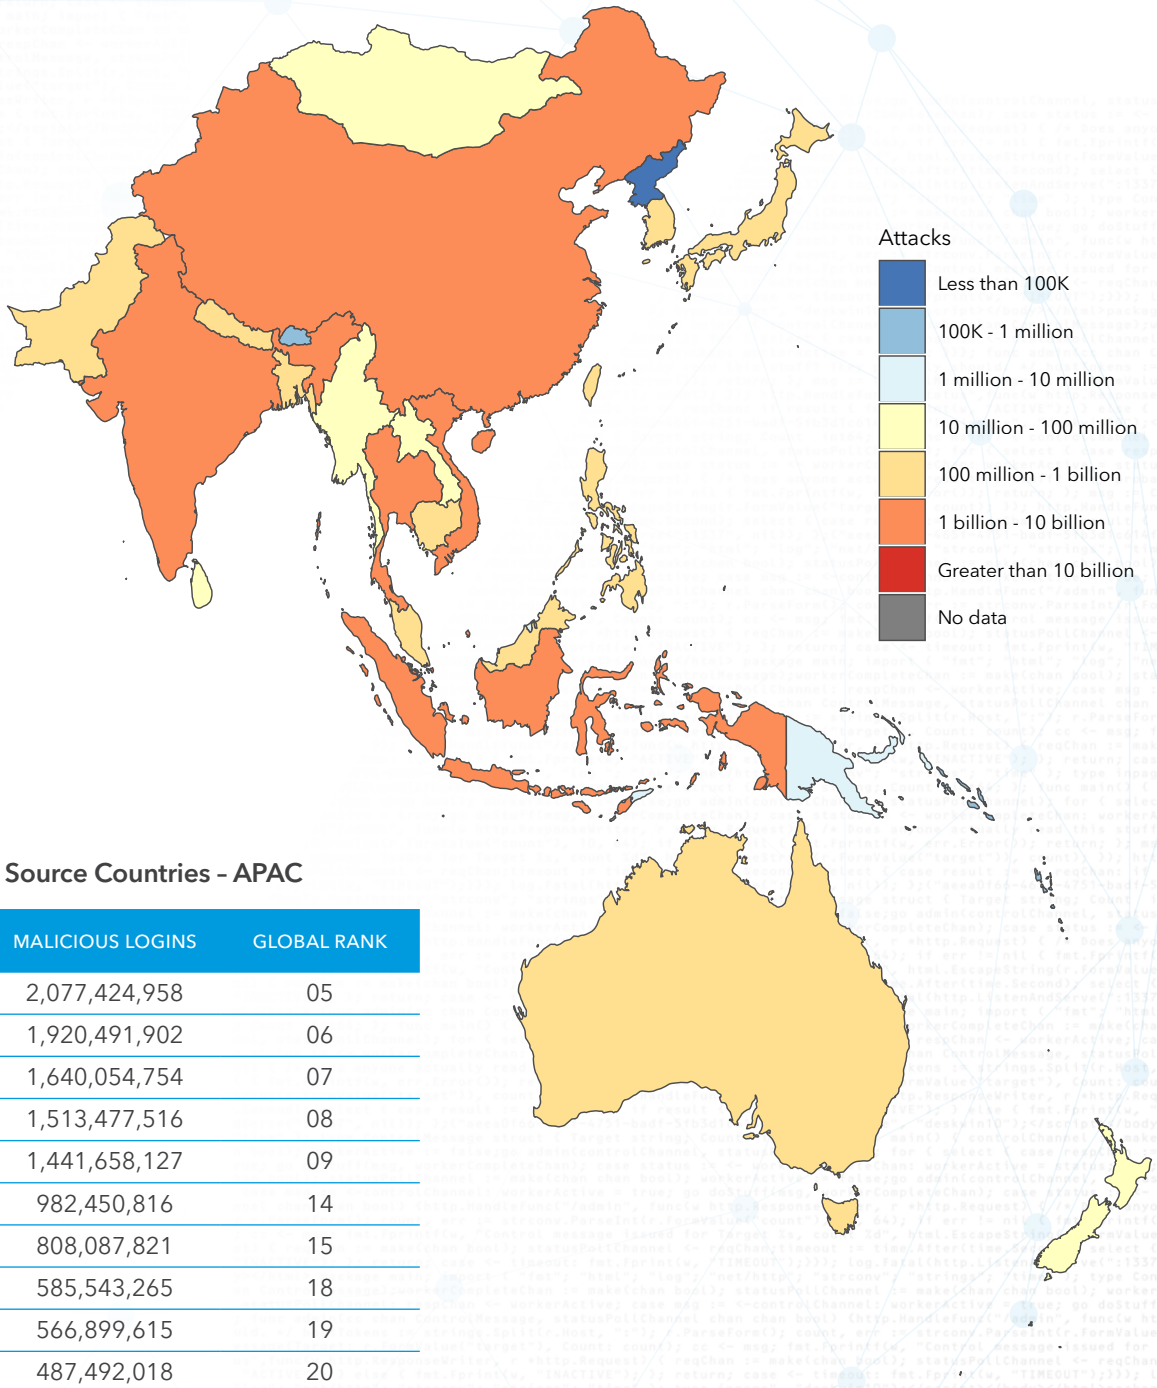

Top 10 Source Countries – APAC

COUNTRY	MALICIOUS LOGINS	GLOBAL RANK
China	2,077,424,958	05
India	1,920,491,902	06
Thailand	1,640,054,754	07
Indonesia	1,513,477,516	08
Vietnam	1,441,658,127	09
Malaysia	982,450,816	14
Singapore	808,087,821	15
Japan	585,543,265	18
South Korea	566,899,615	19
Taiwan	487,492,018	20

Web Application Attack Sources – Americas

November 2017 – April 2019

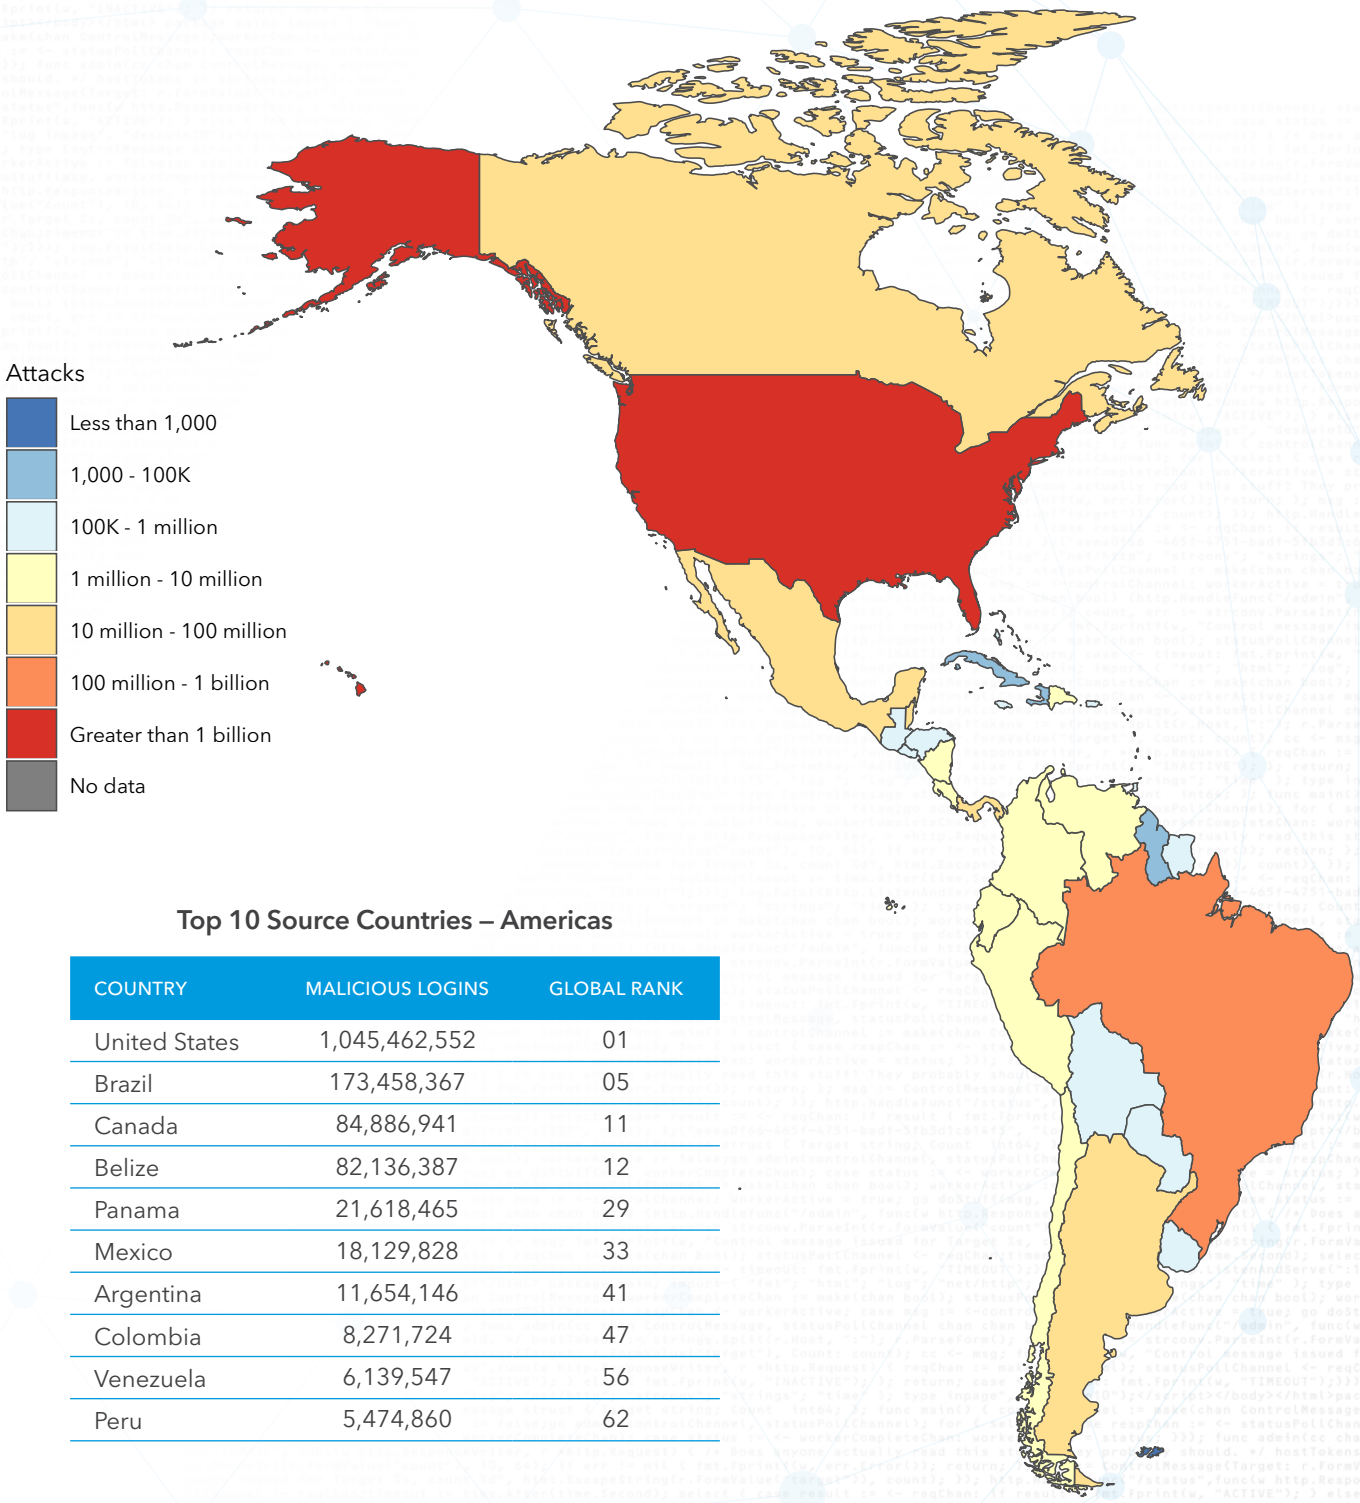

Top 10 Source Countries – Americas

COUNTRY	MALICIOUS LOGINS	GLOBAL RANK
United States	1,045,462,552	01
Brazil	173,458,367	05
Canada	84,886,941	11
Belize	82,136,387	12
Panama	21,618,465	29
Mexico	18,129,828	33
Argentina	11,654,146	41
Colombia	8,271,724	47
Venezuela	6,139,547	56
Peru	5,474,860	62

Web Application Attack Sources – EMEA

November 2017 – April 2019

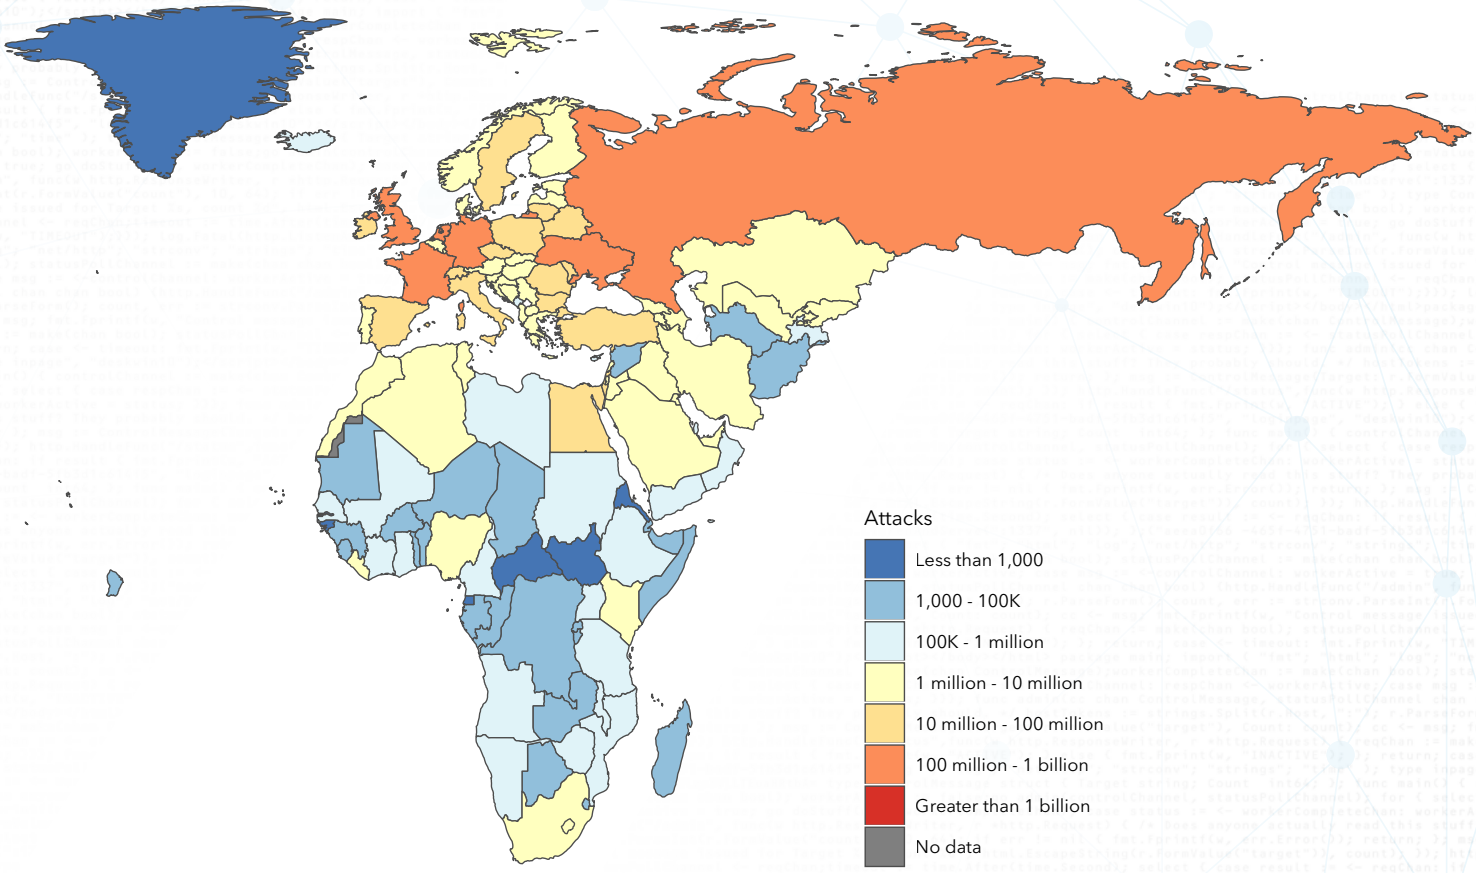

Top 10 Source Countries – EMEA

COUNTRY	MALICIOUS LOGINS	GLOBAL RANK
Russia	696,598,218	02
Netherlands	290,915,336	03
Ukraine	172,452,097	06
France	126,120,125	08
Germany	120,111,620	09
United Kingdom	110,968,477	10
Ireland	74,538,231	13
Turkey	67,507,268	14
Romania	42,416,621	19
Italy	36,639,168	21

Web Application Attack Sources – APAC

November 2017 – April 2019

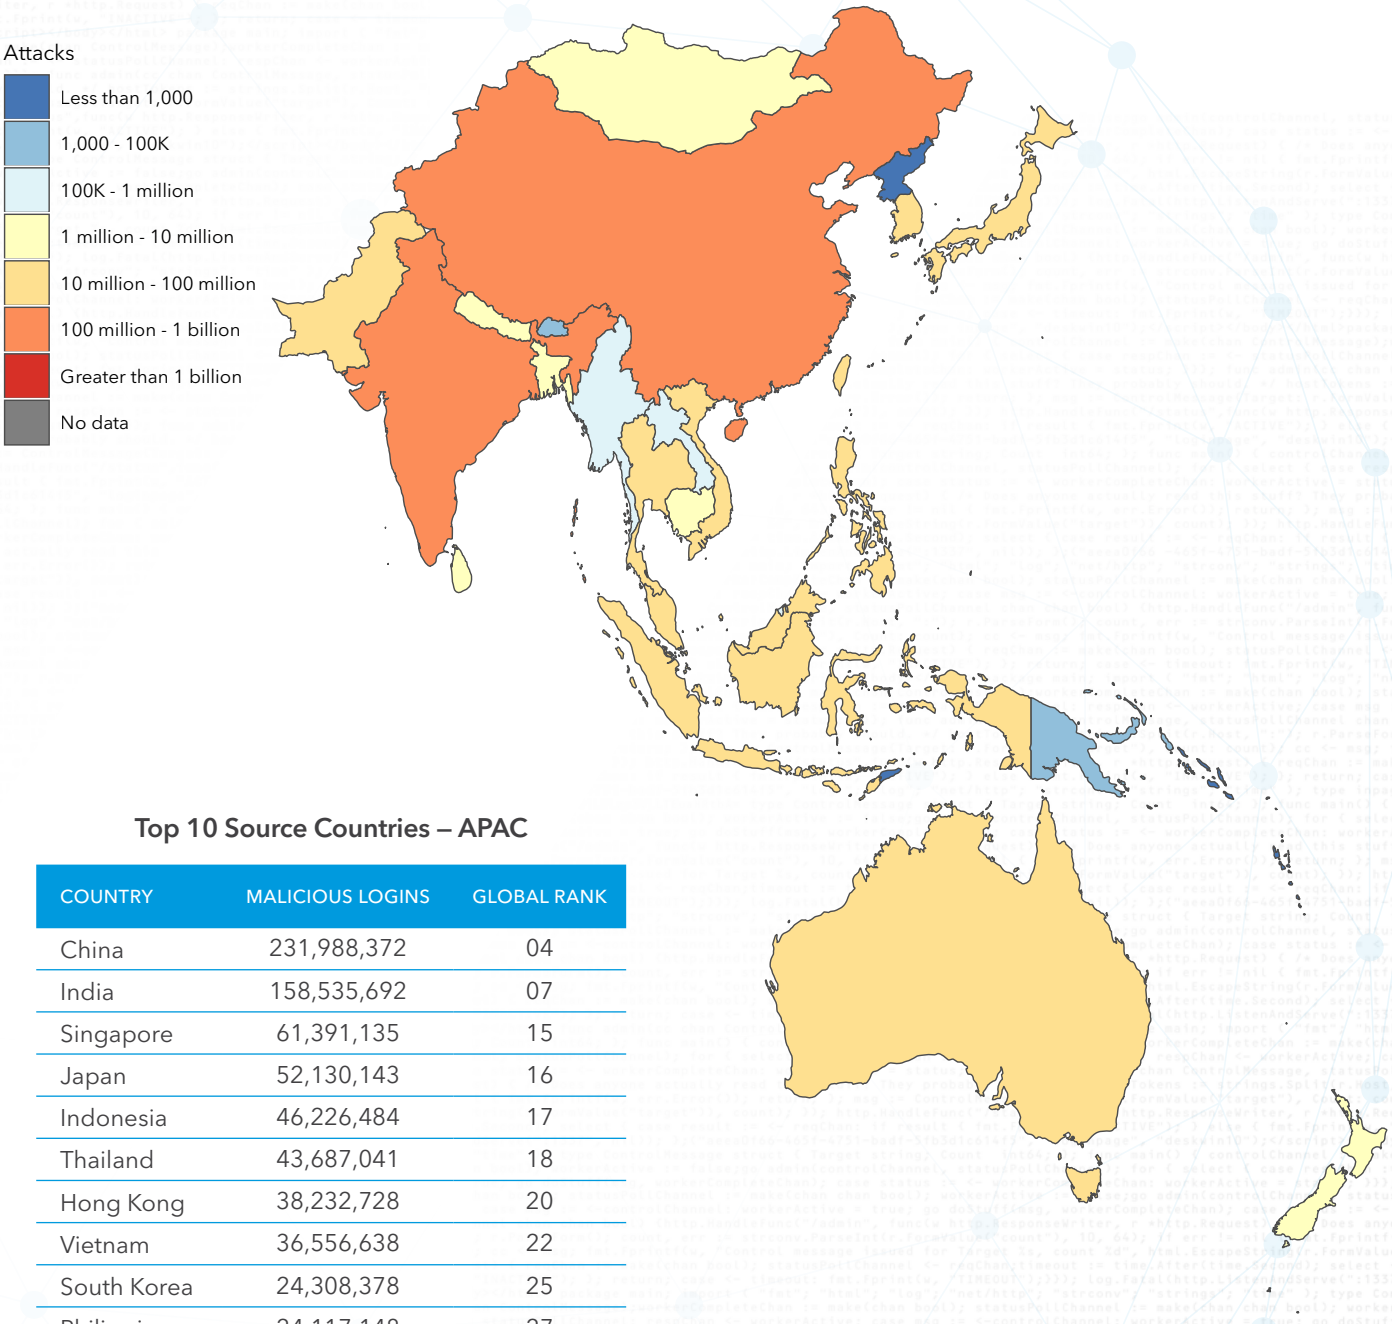

Top 10 Source Countries – APAC

COUNTRY	MALICIOUS LOGINS	GLOBAL RANK
China	231,988,372	04
India	158,535,692	07
Singapore	61,391,135	15
Japan	52,130,143	16
Indonesia	46,226,484	17
Thailand	43,687,041	18
Hong Kong	38,232,728	20
Vietnam	36,556,638	22
South Korea	24,308,378	25
Philippines	24,117,148	27

Top 20 Countries Responsible for Credential Stuffing Attacks

November 2017 – April 2019

COUNTRY	FINSERV MALICIOUS LOGINS	MALICIOUS LOGINS	GLOBAL RANK	PCT FINSERV
United States	1,133,026,190	18,542,844,141	01	6.11%
China	251,719,283	2,077,424,958	05	12.12%
Malaysia	227,443,763	982,450,816	14	23.15%
Brazil	173,176,382	3,197,885,812	03	5.42%
Germany	137,863,141	1,357,470,150	12	10.16%
India	136,925,794	1,920,491,902	06	7.13%
Indonesia	135,309,307	1,513,477,516	08	8.94%
Russia	126,486,202	5,389,826,207	02	2.35%
Vietnam	116,863,450	1,441,658,127	09	8.11%
France	99,191,945	1,358,348,831	11	7.30%
Thailand	80,162,996	1,640,054,754	07	4.89%
South Korea	69,014,099	566,899,615	19	12.17%
Canada	61,212,463	2,378,886,475	04	2.57%
Italy	40,111,943	621,338,014	17	6.46%
Ukraine	39,497,555	781,992,409	16	5.05%
United Kingdom	39,401,694	1,086,213,915	13	3.63%
Japan	37,638,674	585,543,265	18	6.43%
Netherlands	37,438,633	1,361,410,559	10	2.75%
Philippines	33,534,085	218,885,946	35	15.32%
Egypt	30,982,012	292,310,577	30	10.60%

Credits

State of the Internet / Security Contributors

Elad Shuster
Senior Lead Security Researcher – Web Attacks by the Numbers

Or Katz
Principal Lead Security Researcher – Phishing

Lydia LaSeur
Data Scientist – Akamai Research

Steve Ragan
Senior Technical Writer – Akamai Research, Financial Services Attack: A Personal Story

Editorial Staff

Martin McKeay
Editorial Director

Steve Ragan
Senior Technical Writer, Editor

Amanda Fakhreddine
Senior Technical Writer, Managing Editor

Lydia LaSeur
Data Scientist

Marketing

Georgina Morales Hampe
Project Management, Creative

Murali Venukumar
Program Management, Marketing

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 on Twitter. You can find our global contact information at www.akamai.com/locations. Published 07/19.