

WEBROOT[®]
Smarter Cybersecurity[™]

2019 WEBROOT

THREAT REPORT

What's inside

- 4** Webroot Perspective
- 6** Polymorphic Malware and PUAs
- 10** Malicious IP Addresses
- 14** High-Risk URLs
- 18** Phishing Attacks and URLs
- 21** Malicious Mobile Apps
- 22** Summary and Predictions for 2019

THE WEBROOT PLATFORM USES 6TH GENERATION MACHINE LEARNING TO ANALYZE 500 BILLION DATA OBJECTS EVERY DAY.



750M+

Domains



32B+

URLs



4B+

IP Addresses



31B+

File Behaviors



62M+

Mobile Apps



67M+

Connected Sensors

POLYMORPHIC MALWARE AND PUAS



In 2017, 93% of malware and 95% of potentially unwanted applications (PUAs) were polymorphic, and this trend continued throughout 2018. Polymorphic code makes a change to a single instance of malware (through names, encryption keys, signatures, hashes, function instructions or the order of execution flow) so it can be delivered to a large number of people while still evading detection. Because polymorphic malware and PUAs never have the same identifiers, existing signatures will never match the new variant; this means that pattern-matching security products cannot detect new variants quickly enough to prevent infections.

In 2018, 93% of malware seen by Webroot was polymorphic.

Endpoints running Webroot® protection see more than five hundred million brand new, never-before-seen portable executable (PE) files each year, and this number continues to increase. But, while the number itself is going up, the percentage of files that are determined to be malware or PUAs is going down. Of the new PEs seen in 2018, less than 1% were deemed malware, compared to 1.5% in 2017 and 2.5% in 2016. The decline in PUAs in 2018 was even more dramatic (see Figure 1), dropping to 0.11% from 0.4% in 2017 and 2.2% in 2016. We will explore possible reasons for this decline later in this report.

Of the endpoints reporting an infection, 68% were consumer devices, while 32% were business endpoints. When we look at the average number of malware files per device, we see a dramatic decline in 2018 (Figure 2).

	2016	2017	2018
Percent of executable files that are malware	2.5%	1.5%	0.88%
Percent that are PUAs	2.2%	0.4%	0.11%

Figure 1: PEs determined to be malware or PUAs

	2016	2017	2018
Total malware files per device	0.66	0.48	0.07
Malware per consumer device	0.59	0.53	0.09
Malware per business device	0.61	0.42	0.04

Figure 2: Average number of malware files per device

In Figure 2, we see that, on average, consumer devices are infected more than twice as often as their business counterparts. However, the business landscape is not populated solely by corporate-owned PCs. Many companies allow their employees to connect their personal devices, including PCs, to the corporate network, which greatly increases the level of risk to the organization.



While the decline in malware is real, it is hardly mission complete. We've seen continued innovation in tactics and techniques, particularly with Emotet and Trickbot in 2018. Adding UPnP and Tor functionalities, respectfully, have made these threats more resilient and difficult to knock offline.

GRAYSON MILBOURNE | SECURITY INTELLIGENCE DIRECTOR



THE ROLE OF OPERATING SYSTEMS

The operating system (OS) plays an important part in the decrease in new malware and PUA files seen. As we saw last year, the move to Windows® 10, which is a generally safer OS in which antivirus is always on, helps explain the downward trend.

Devices that use Windows 10 are at least twice as secure as those running Windows 7.

Of business endpoints running Webroot protection, more use Windows 10 than Windows 7 (45% and 43%, respectively), while just 3% run Windows 8 and only 1% still use Windows XP (which Microsoft stopped supporting several years ago). In fact, business use of Windows 10 reached a tipping point in November 2017 and has been steadily increasing at a rate of about 1.2% per month. Nevertheless, growth is slower than one might expect, given the obvious security benefits of this newer operating system, and we anticipate that it will be two or three more years before we see Windows 10 usage in the

business world on a par with consumer adoption. In general, the business sector has been slower to move away from older versions of the Windows operating system, perhaps due to software requirements for legacy operating systems like Windows 7 and XP. On the consumer side, Windows 10 adoption remains steady at approximately 75%, without much movement throughout the year. Meanwhile, Windows 7 stands at 13% and Windows 8 at 9%.

Overall, Windows 7 shows a higher rate of infections per endpoint device (.07) than Windows 10 (.05). However, when looking at consumer versus business devices, the story is quite different. Consumer systems saw more than twice as many infections per endpoint (.09) as business systems (0.04 per endpoint). The numbers are even more striking when viewed in terms of the OS. Consumer systems running Windows 7 saw an average of 0.18 infections, while business systems running that version of the OS saw only an average of .04. For Windows 10, consumer endpoints saw an average of .07 infections, whereas business endpoints saw, on average, only .02. Almost all represent decreases from previous years (see Figure 3).

		2016	2017	2018
Consumer	Average	0.66	0.48	0.09
	Windows 7	0.59	0.53	0.18
	Windows 10	0.61	0.42	0.07
Business	Average	0.11	0.07	0.04
	Windows 7	0.19	0.07	0.04
	Windows 10	0.05	0.03	0.02

Figure 3: Infections per endpoint

Over the last year, we have seen a relatively steady decline in malware on Windows 10 machines for both consumer and business. There were spikes for malware on business PCs running Windows 10 in the period from September through December, likely due to back-to-school and holiday malware campaigns.

While malware and PUAs are decreasing as a percentage of new files seen, the numbers are still noteworthy, and the threat is not going away. There are several possible reasons for the decrease in the number of malware and PUA files seen by Webroot-protected devices.

- » First, Webroot now detects malicious activity even earlier in the kill chain. In other words, we block executable files from infecting endpoints via malicious URLs (see URL section on page 14) and prevent executables on endpoints from downloading additional bad executables.
- » Second, changes in the malware ecosystem—e.g. the fact that drive-by download exploits have largely dried up and attackers are finding easier ways to profit than installing malware—contribute to the reduction in the percentage of new files classified as malicious.
- » Third, as we have mentioned, Windows 10 is a safer operating system than others. Windows Defender® will activate itself on a system if other antivirus solutions become inactive. Because of this, security is stronger overall and there are fewer infections.

When it comes to PUAs, companies like AppEsteem¹ may also contribute to the decrease: they enable end users to download and use apps with less risk. App vendors can develop and deliver clean apps by following clear rules that are reviewed by cybersecurity companies, thereby being certified as safe.

A DEEPER DIVE INTO MALWARE

Webroot has added more sources of information and analysis during 2018, and the more information we can provide to administrators, the better they can protect their organizations. This year, new data includes the number of machines on which malware is seen and the locations where threat actors hide malware.

Over half of devices that became infected once were re-infected within the same year.

In 2018, we found that 93% of malware was only seen on a single PC, and of the machines that were infected, over half (54%) saw more than one infection over the course of the year. More than 39% of consumer endpoints that were infected at least once saw between 2-5 infections in 2018, while the percentage for business endpoints was slightly lower, at 35%.

This is likely the result of multiple polymorphic files attacking individual devices. Additionally, when Webroot protection is first installed on a PC, it often finds multiple current infections. Another reason for the high incidence of PCs reporting more than one infection is that some malware drops multiple files (see Swiss Army Knife section on page 19). The net result is that administrators must remain vigilant; endpoints infected once are likely to become infected again over the course of a year.

Nearly 54% of malware in 2018 hid in the %appdata% and %temp% folders.



Repeat infections often have behavior-based causes. Users who frequent torrent sites for game cheats or activation keys should know these are almost always infected. Other behaviors like installing untrustworthy apps without carefully evading bundled software also contributes to repeat infections.

GRAYSON MILBOURNE | SECURITY INTELLIGENCE DIRECTOR



New in this year's report is detailed information on where malware tries to install itself. We compile this data by normalizing the paths from various Windows operating systems and versions to display a consistent view of the paths, and then base the percentages on new file encounters and ratios of where malware is found. Based on our analysis, we have found several likely spots where malware hides, including %appdata%, %temp%, and others—although, realistically, malware can hide almost anywhere. See Figure 4 for details.

Malware Installation Locations	%
%appdata%	29.4%
%temp%	24.5%
%cache%	17.5%
%windir%	12.3%
%desktop%	6.1%
%programfiles%	4.4%
OTHER	5.8%

Figure 4: Common malware installation locations on Windows machines

The %appdata% directory is a good example of why location matters. Malware authors often try to install the main launching application into a subdirectory of the %appdata% folder, which contains application settings, files, and data specific to the apps on a Windows PC. For consumers, we detected 31% of the %appdata% files as bad, whereas only 24% of the %appdata% files were found to be malicious for business users. The higher percentage of malicious files found in consumer folders is likely because %appdata% is only updated when a new app is installed, and app churn is lower in the consumer space than business. However, when it happens, these apps are more likely to be infections, and a consumer PC is typically less secure than a business one.

In looking at new files seen in 2018, other folders where malware often hides include:

- » **%temp%** – 24.5% overall, 30% for business, 23% consumer
- » **%cache%** – 17.5% overall, 11.3% for business, 19.2% for consumer
- » **%windir%** - 12.3% overall, 19.2% for business, 10.5% for consumer

We encourage administrators to create policies that look for anomalous behavior in the %cache%, and %temp% directories; for example, policies could prevent files from executing from these directories. By doing so, admins could eliminate more than 40% of the opportunities malware has to launch itself.

Overall, we can conclude that while there are somewhat fewer instances of malware and PUAs, the problem hasn't gone away; it has just changed, and extreme vigilance is needed. Efforts to make operating systems cleaner and render it more difficult to install PUAs have contributed to an overall decrease. We see that business PCs are safer than consumer PCs, and those running Windows 10 are safer than those running Windows 7. However, the trend of allowing employees to bring their personal devices to work and connect them to the corporate network increases organizational risk.

WHAT HAPPENED TO RANSOMWARE?



While the ransomware attacks of 2017 spread fear and panic across the globe, with companies scrambling to safeguard mission-critical data and paying millions in ransoms via cryptocurrency, the story in 2018 was much different. Ransomware has proven to be an effective tool to extract money from targets who are not prepared, but potential victims are now doing a better job of securely backing up their data, making it more difficult for attackers to scare users into paying the ransom. As a result, ransomware has evolved to be more targeted, better implemented and, therefore, much more ruthless.

Some examples of ransomware attacks in 2018 include attacks on ports in Barcelona, San Diego, and Long beach; airports in Bristol and Atlanta were hit with ransomware attacks, and government and health care organizations (including one US hospital that paid a \$55K ransomⁱⁱ) were targeted this year. One very high-profile ransomware attack, known as SamSamⁱⁱⁱ, encrypted hundreds of networks in the U.S. and other countries and resulted in more than \$6 million in payouts before the criminals were indicted. Total damages from that attack on more than 200 victims exceeded \$30 million. SamSam used the Remote Desktop Protocol (RDP) as an attack vector, leveraging tools such as Shodan to scan for systems with inadequate RDP settings. These unsecured RDP connections may be used to gain access to a given system and browse all its data as well as shared drives, providing criminals enough intel to decide whether to deploy ransomware or some other type of malware. Despite repeated warnings, many companies still do not protect their RDP connections sufficiently.

MALICIOUS IP ADDRESSES



Each year, Webroot sees millions of malicious IP addresses which are used to send out spam, distribute malware, obfuscate the origin of malicious traffic, or otherwise enable bad actors to wreak havoc with consumer and business computers. We track these IP addresses by the malicious activities they carry out: scanners or proxies, spam, Windows exploits, web or denial of service (DoS) attacks, botnets, phishing, and mobile threats. The best way to address the potential danger of malicious IPs is to block them automatically so they cannot do damage. But doing this requires an in-depth understanding of IPs, their locations, and their actions, in order to proactively block them.

Of the malicious IP addresses we saw in 2018, a colossal 82% were categorized as malicious because of spam (see Figure 5). The figure itself represents a significant increase from previous years: in 2017, 65% of the malicious IPs found were sending spam.

Beyond spam, the most frequent activity seen in relation to malicious IPs was open proxies at 9% (including anonymous and Tor, which allow anonymous traffic to pass through), followed by botnets at 4.3%, up from last year's 3% rate.

Scanners came in at only 2% but still represent a troublesome threat; hackers scan environments to learn specifics about the network configuration, software installed, and user data so they can tailor attacks to that particular environment. Windows exploits dropped dramatically from last year's 9% to just 1.1%. While the move to Windows 10 factors into this drop, another important reason is activity in the white hat space, e.g. Google Project Zero, which results in earlier discovery of vulnerabilities, greater disclosure, and fewer exploit kits. Windows exploits remain a popular method for distributing malware because they leverage a vulnerability in the operating system, software,

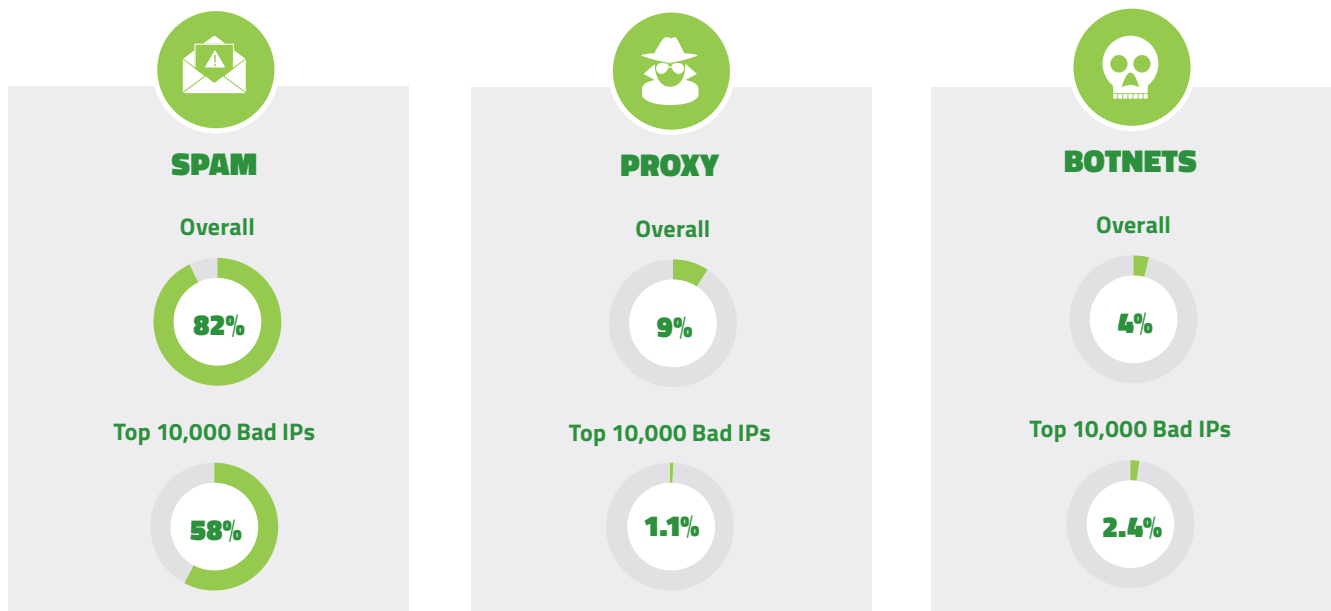


Figure 5: Categories of malicious IP addresses in 2018



Exploit kits are still prevalent but are much less effective against fully patched systems. The impact of bug bounty programs, Google’s Project Zero and automated fuzzing tools have greatly improved both the security and stability of software making it more difficult to exploit as a vector for infection.

GRAYSON MILBOURNE | SECURITY INTELLIGENCE DIRECTOR



25 of the top recurring malicious IP addresses were reused 95 times.

browser, or plug-in; consequently, we expect the activity here to ebb and flow as threat actors save exploits for high-profile targeted attacks.

It’s important to keep in mind that IP addresses are not static and may cycle from malicious to benign and back multiple times. While 60% of the millions of malicious IP addresses we saw in 2018 only appeared on the list once, hundreds of thousands appeared at least two or more times. In fact, 25 of the top recurring IP addresses were reused 95 times. IP addresses are often reused to carry out botnet, spam, and scanner activity. Another review showed that 193 IPs accounted for more than 11,000 transitions (that is, IPs that were used for botnets, then

for benign purposes, then again for botnets multiple times during the year.) For spam, 8,704 IPs accounted for more than 185,000 transitions, and 4,848 scanner IPs were used in more than 103,000 transitions.

Here’s an example of how this works: we see a server sending out spam emails and we identify the IP address as a spam distribution node. Because of this undesirable behavior, we add the IP to a blacklist. But it doesn’t stay there indefinitely. IPs on the blacklist are revisited to see if they still exhibit malicious behavior. If not, they leave the blacklist, but we record the historical behavior of each IP address, which influences their IP reputation scores. Hundreds of thousands of new IPs are added to and removed from the blacklist multiple times a day.

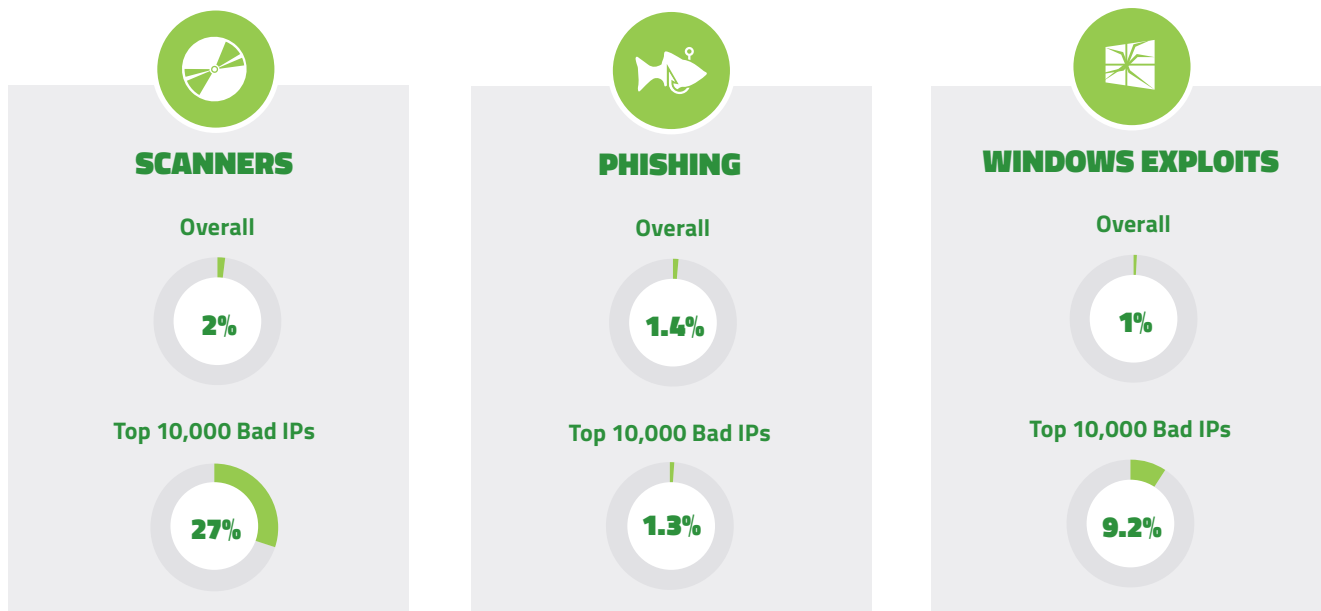


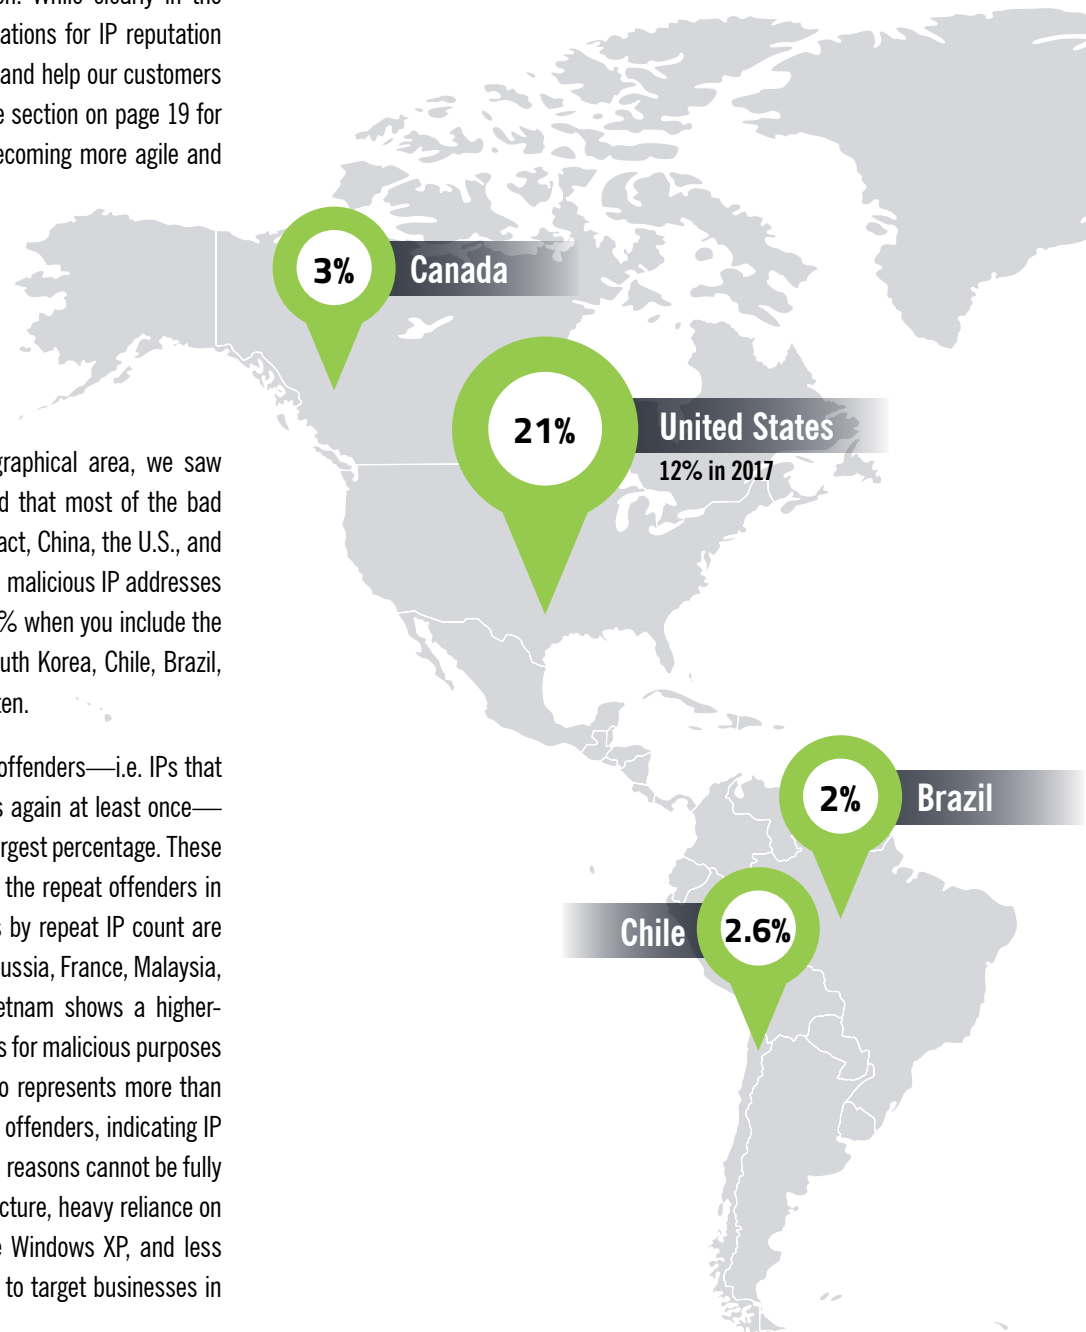
Figure 5: Categories of malicious IP addresses in 2018

We often find the same IPs reused for malicious purposes on a later date. Of the top 10,000 IPs we analyzed, 48.6% of them were recorded as malicious 20 times or more. More than 150 of them traversed the black list 50 or more times, and a small percentage (1.5%) were on and off the black list even more often than that.

Interestingly, we found that, of the top 10,000 most recurrent IP addresses, more than half were used for multiple types of malicious activity. For example, they might be used for web attacks, Windows exploits, and proxy or Tor activity. We found that 16 addresses were used for six different major types of malicious activity, and two were used for all seven. While clearly in the minority, this multi-use tactic has implications for IP reputation services and informs how we track them and help our customers deal with them. See the Swiss Army Knife section on page 19 for an in-depth look at how attackers are becoming more agile and using multiple tools to carry out attacks.

When we look at malicious IPs by geographical area, we saw activity from all 195 countries, but found that most of the bad IPs came from a handful of locations. In fact, China, the U.S., and Vietnam account for more than 60% of all malicious IP addresses seen. The figure goes up to more than 80% when you include the rest of the top ten: Germany, Canada, South Korea, Chile, Brazil, Russia and India. Figure 6 shows the top ten.

When we focus on the top 10,000 repeat offenders—i.e. IPs that are malicious, benign, and then malicious again at least once—we see that Vietnam and China have the largest percentage. These two countries account for almost 57% of the repeat offenders in the top 10,000 list. The top ten countries by repeat IP count are Vietnam, China, the U.S., Hungary, India, Russia, France, Malaysia, Brazil, and Switzerland. Interestingly, Vietnam shows a higher-than-normal hosting/reuse of IP addresses for malicious purposes with 662,000 unique IP addresses. It also represents more than 40% of the top 10,000 in terms of repeat offenders, indicating IP reuse is rampant in Vietnam. Although the reasons cannot be fully determined, it is likely that aging infrastructure, heavy reliance on older, less-secure operating systems like Windows XP, and less stringent law enforcement make it easier to target businesses in Vietnam and exploit their networks.





China, the U.S. and Vietnam account for 60+% of all malicious IPs seen.

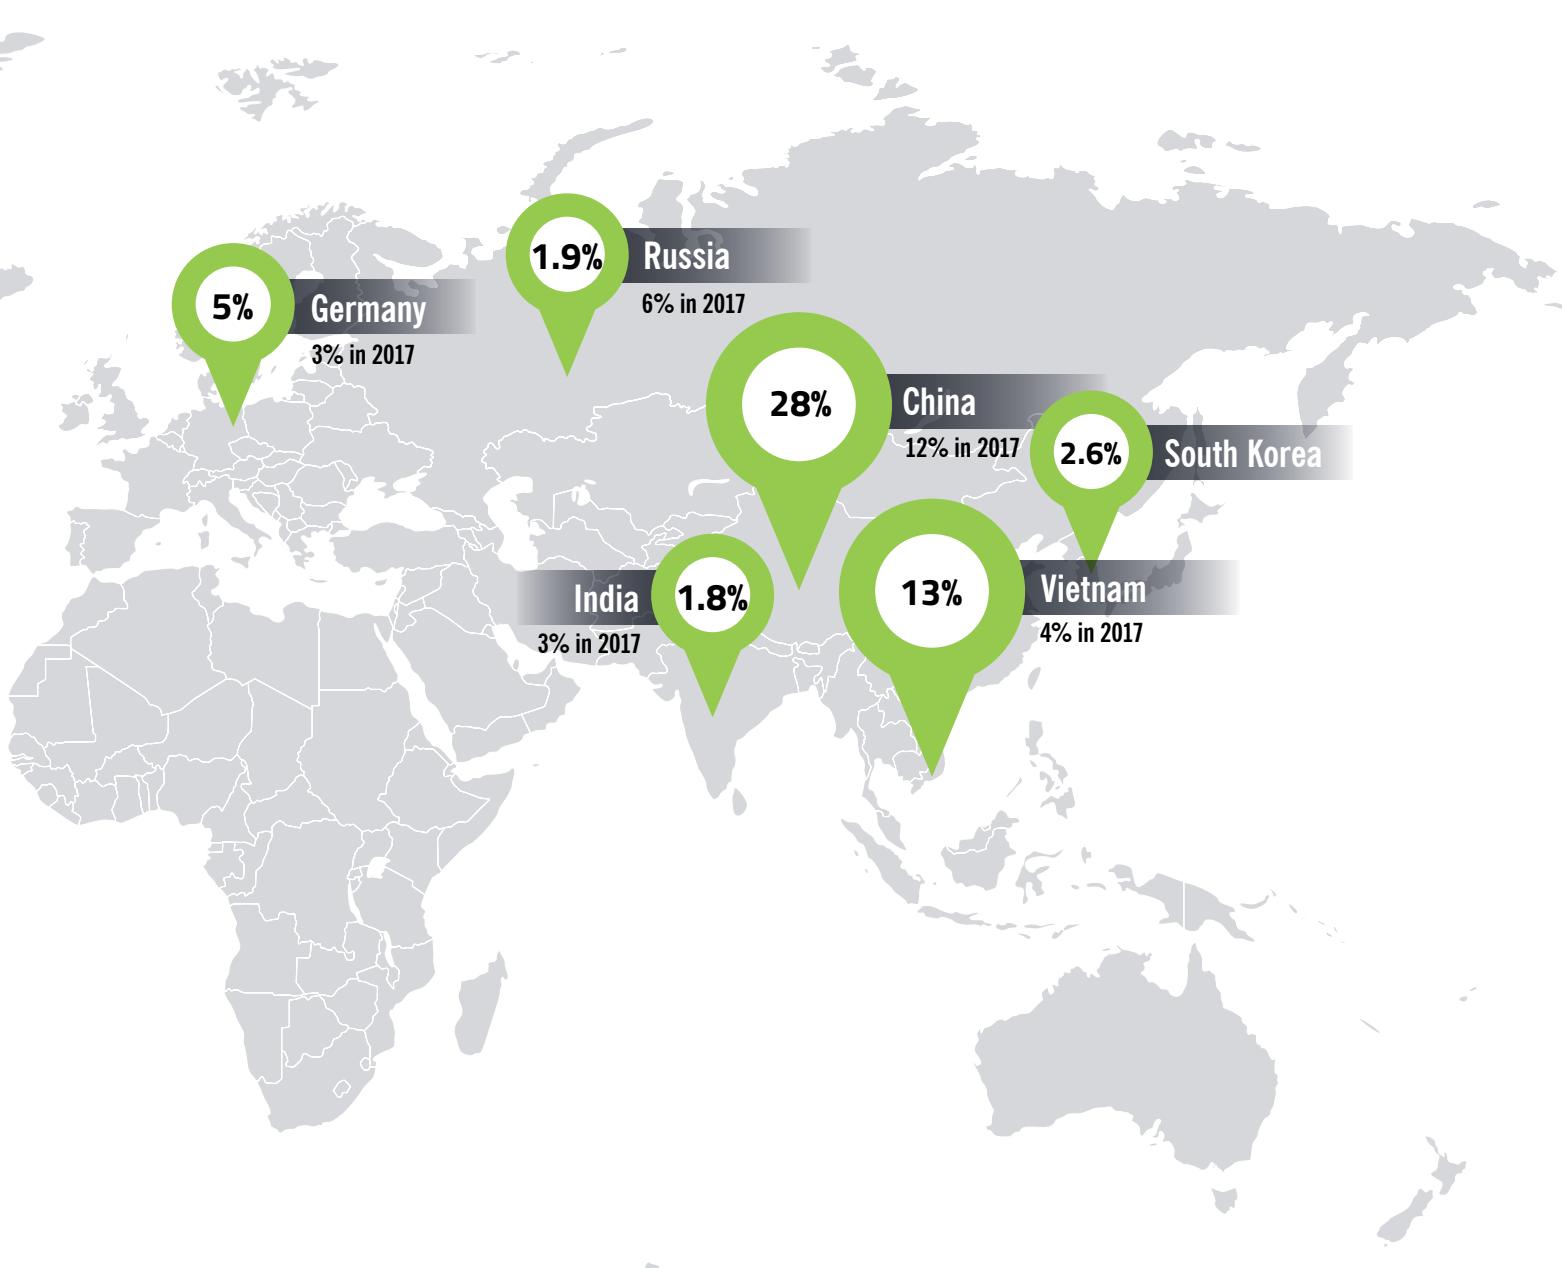


Figure 6: Top 10 countries for malicious IP activity

HIGH-RISK URLs



Webroot has categorized more than 32 billion URLs to date, continuously examining their history, age, popularity, location, networks, links, real-time performance, and behavior. We classify them in terms of their primary purpose (e.g. shopping, adult sites, gambling, etc.) or malicious intent (such as phishing, botnets, malware sites, spam sites, etc.) This intelligence is provided to network and security vendors via Webroot BrightCloud® Web Classification and Web Reputation Services, which organizations can use to set web policies to protect users.

Between January and December 2018, the number of phishing sites detected grew 220%.

In 2018, Webroot classified hundreds of millions of unique URLs, approximately 3% of which were classified as high risk. The number of phishing sites per day, in particular, more than doubled over the course of the year. The number of botnets fluctuated, but there were spikes in October and November, likely related to virulent attacks such as Emotet (described in the Swiss Army Knife section on page 19.) Some fluctuations in the level of activity are due to seasonal factors, such as the fact that attacks ramp up during back-to-school time and the holidays.

Malware is undoubtedly the most interesting category of high-risk URLs, due to the prevalence of cryptojacking and cryptomining activities. These sites, which leverage site visitors' hardware to mine cryptocurrency represent a significant percentage of the malicious URLs seen by Webroot, and recent enhancements in tracking will undoubtedly reveal far more than we have seen to date.

CRYPTOJACKING AND CRYPTOMINING

Cryptojacking is the practice of using browser-based programs that run via scripts embedded in web content to mine cryptocurrency with unused CPU without the user's knowledge or consent. Cryptomining is when criminals install malware that usurps a user's CPU to mine cryptocurrency. Both of these techniques have grown rapidly to become major threats, as they can be more lucrative than ransomware attacks, and have a smaller illegal footprint. Monero continues to be the cryptocurrency of choice for these types of attacks for a variety of reasons. Its innovative use of ring signatures and decoys to hide the origin of the transaction ensures that it's completely untraceable. Once criminals receive payment to a Monero wallet address, it can be sent to an exchange where it can be cashed it out as if it were clean.



Cryptojacking has been out over a year now, so people are aware of it and know how to block it. But the way they're blocking it with browser add-ons is very rudimentary. As this threat evolves and criminals start to obfuscate the domains, those add-ons will become obsolete and real-time threat intelligence will be the only effective way to block cryptojackers.

TYLER MOFFITT | SECURITY ANALYST



In 2018, millions of URLs used cryptojacking.

Beyond the obvious utility of an untraceable transaction system that doesn't need to be laundered, Monero is one of the few cryptocurrencies that maintain ASIC resistance. Most cryptocurrencies use a proof of work mining system, but the algorithm used to mine the blocks can be worked by a specialized chip (ASIC), which is designed to specifically hash that algorithm and mine that cryptocurrency. Companies like Bitmain have thrived on creating ASIC machines that are so much more efficient they render consumer-grade CPUs and GPUs obsolete. The Monero development team maintains an algorithm that makes sure it stays ASIC-resistant. This means that Monero can be mined profitably using consumer-grade CPUs and GPUs in personal computers and will be able to do so for the foreseeable future. Massive price fluctuations have caused some ups and downs in the popularity of both types of attack, as the entire crypto market has lost 85-90% of its value over the past year.

Nevertheless, huge robberies, hacks and mining operations made the news in 2018:

- » Strong arm robberies of high-profile crypto investors and traders resulted in the loss of assets in the millions of dollars, and sometimes even loss of life.^{iv}
- » More than \$731 million worth of cryptocurrencies were stolen from crypto exchanges in the first half of 2018.^v

Despite the decrease in cryptocurrency prices, the number of sites using these attack methods has increased significantly over 2018. Since the victim is responsible for the power bill and the increased stress on their hardware associated with illicit mining, these attacks cost next to nothing and start making money for criminals almost immediately. As long as cryptocurrency is worth something, the potential profits will still be very attractive for criminals. The level of cryptojacking URLs we saw each month in the first half of the year more than doubled in the period from September through December. Cryptojacking will remain popular, as it generates profits without needing to infect a PC. We have seen increased competition in web-based cryptojacking with several new sites offering the functionality in 2018. For a long time, the only real player was Coinhive, who created this new type of attack back in 2017. While Coinhive still dominates with more than 80% market share, some new copycat cryptojacking scripts are gaining in popularity, such as Cryptoloot, JSEcoin, Deepminer, Coinimp, Minr, and Omine.

The detection rate for cryptojacking by Webroot end users attempting to visit those URLs showed a steady decline in 2018, despite a spike in September (see Figure 7). Of the annual detections, 11% were found in January, compared to just 5% in December. This decrease can likely be attributed to the decline in cryptocurrency value coupled with improved detection and blocking capabilities at the network layer.




Figure 7: Cryptojacking detections throughout 2018

We have also seen evidence of browsers improving their protection capabilities against cryptojacking during that time. However, we anticipate a future decline in the effectiveness of cryptojacking prevention from browsers and extensions, as criminals experiment with obfuscation to avoid detection.

In some ways, mining is a more flexible and profitable method than cryptojacking, since the damage does not cease if the user navigates away from the hijacked site, and is not easy to curtail via JavaScript restrictions. Rather, cryptomining relies on persistence and obfuscation, and attempts to hide the fact that it increases CPU load by running while the CPU is idle. More than 4.3% of all Monero in circulation has been mined illicitly through victims' hardware.^{vi} It is also worth noting that botnets (see the Swiss Army Knife section on page 19) dropped cryptominer payloads in 2018, with XMRig being the favorite. Mining presents a path to profitability when no mission-critical data is found, and when ransomware isn't a suitable option, all without involving user consent or willingness to pay a ransom.

MALICIOUS URLS BY COUNTRY

By far the biggest contributor of malicious URLs was the U.S. at 63%, followed by China at 5% and Germany and Hong Kong at 4% each. Russia and the Netherlands each came in at 3%.

The fact that the majority of sites hosting malware are located in the U.S. is a big change compared to previous years, in which the U.S. accounted for just 12% in 2017 and 22% in 2016. The percentage may have grown so drastically because so many sites in the U.S. were infected with, or used, cryptojacking.

URL OBFUSCATION

When users click a button or hyperlink, it is sometimes difficult to know where the web browser will actually take them. Sometimes, this occurs without malicious intent, such as when a user is asked to "click here" to get more information about an organization, product, or service. However, URL obfuscation is frequently used to direct users to a malicious site, or to malicious content on a benign site. A massive 40% of malicious

In 2018, just ten countries hosted the majority of high-risk URLs.



Figure 8: Regions with the most high-risk URLs

A massive 40% of malicious URLs were found on good domains.



URLs were found on good domains, since legitimate websites are frequently compromised to host malicious content. Those who use intermediary devices without SSL inspection capabilities should be aware of potential loopholes in their security policies due to this behavior.

Two of the most common ways to hide the true destination of a URL are through URL shorteners and cloud storage. URL shorteners are very popular, easy to use, and convenient for Twitter and other applications that restrict content posts to a handful of characters. The danger of shortened URLs is that they do not disclose the actual destination. They also pose risks for users and servers, as they can be used to circumvent SSL and carry out browser attacks, one-click exploits, and DoS attacks.^{vii} At Webroot, we evaluated thousands of URL shorteners and found that they can represent a significant risk to users. Examples include bit.ly and goo.gl. We found that users had a 1 in 130 chance of clicking on a bit.ly URL and going somewhere malicious, and a 1 in 190 chance with goo.gl, which was mostly discontinued by Google in May 2018 (with small exceptions, such as sharing in Google™ Maps). Despite efforts by Google and others to make sure shortened URLs are valid, we still see thousands that are malicious.

A similar opportunity to hide the URL comes when data is stored in the cloud. The user is given a link on a domain, which might itself be benign. However, when we looked at the path of such URLs, we found that at least one-quarter contained malicious content. To the casual observer, or to tools that only see the domain name, the link looks fine, since the domain is known and trusted. However, trusted sites may be compromised—even if only for a short period of time before being discovered—and threat actors know this is an effective method for evading detection. As an example, if a Photobucket account is compromised, it may be used to host malicious files for distribution.

This underscores the importance of leveraging URL-level intelligence, and not simply relying on URL whitelisting. Webroot evaluates the entire path, to ensure that users are not exposed to malicious URLs.

PHISHING ATTACKS AND URLS



Compared to 2017, 2018 saw a 36% increase overall in the number of phishing attacks.

As in years past, phishing attacks continue to be one of the most popular ways to attack both businesses and consumers. The number of phishing sites detected grew 220% between January and December of 2018. Compared to 2017, that means 2018 saw a 36% increase overall in the number of phishing attacks. Most phishing attacks impersonate financial institutions; of all the organizations phished, financial institutions represented 77%. Interestingly, in 2018 we found 20 different industries targeted by phishing campaigns. Behind financial institutions, yet still representing a significant percentage, were payment services (4.78%), cryptocurrency and retail (both at 2.53%), and government agencies (2.25%). Other targets included social media, streaming, SaaS, healthcare, delivery services, gaming, file hosting, education, insurance and others.



I would expect cryptocurrencies to occupy a larger percentage of the targets for phishing attempts moving forward, since there are many more targets using exchange or web wallets, which are highly vulnerable to human mistakes.

TYLER MOFFITT
SECURITY ANALYST



While the sites targeted for impersonation vary from year to year, we saw many of the same names in 2018 that were targeted in 2017. Google, Microsoft, Dropbox, Paypal, Apple, Facebook, and Yahoo continue to hold spots in the top ten, joined this year by Adobe and Bank of America (see Figure 9).

Dropbox, in particular, makes an interesting target. A phishing attack that yields Dropbox credentials could reveal sensitive consumer and business data stored on Dropbox, such as financial accounts and other personal information. It could also be used to host malware. If a criminal were to break into a business administrator's account, they could see everything: corporate intellectual property and even cryptokeys, potentially unlocking a massive amount of highly sensitive, mission-critical data.

When we expand the view to the top 20 companies targeted for phishing attacks, we see some new names including Netflix, Target, and Amazon. These are not the usual suspects; it would not be immediately obvious what benefit someone might get from impersonating Netflix. However, because users continue to reuse passwords, targeting less obvious account types might yield credentials that could be used to break into banking or other more lucrative sites. Criminals recognize that all brands have a web presence, and if users continue to use a single password to log into multiple sites, a simple phishing attack could yield tremendous benefits. But phishing attacks targeting Amazon pose threats beyond simply getting access to an individual's Prime account. Gaining access to the master keys to an Amazon Web Services (AWS) account could potentially allow a threat actor to see a company's entire cloud infrastructure, with devastating consequences.

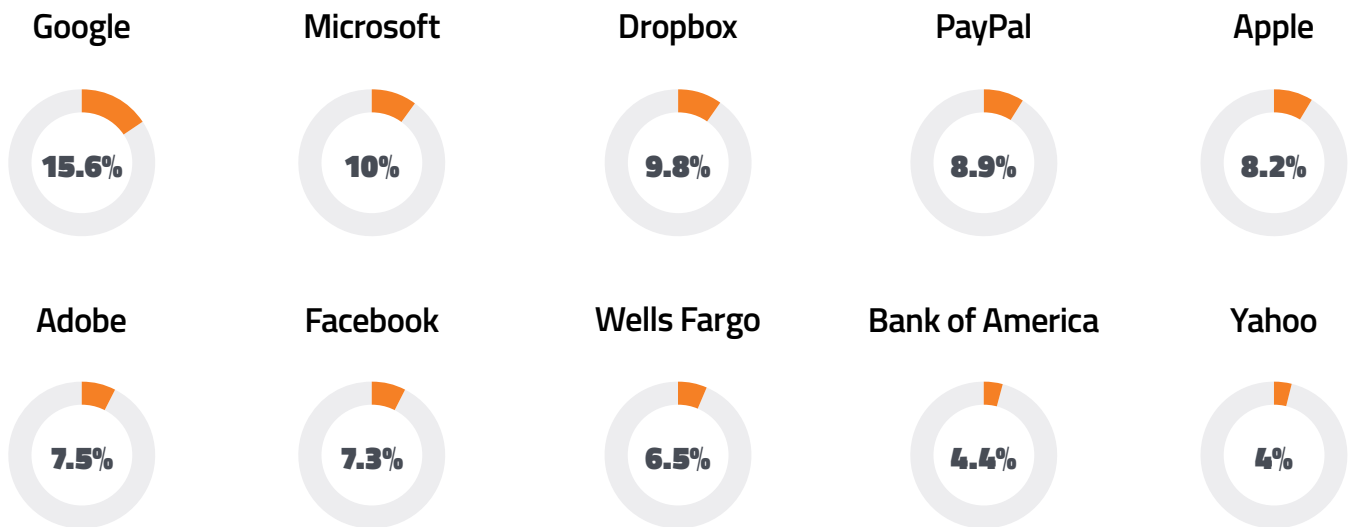


Figure 9: Top 10 Impersonated Organizations

HTTPS HELPS SECURE USERS, BUT THERE ARE CAVEATS

HTTPS was created to protect users by ensuring their sensitive information is encrypted as it travels between the browser and a server, preventing eavesdropping. Users have learned to look for the lock in the browser bar as an indication that a given site is trusted and safe. As with all new security innovations, phishers have found a way to take advantage. They can add SSL certificates to malicious websites, making them appear safe. It is easy to obtain SSL certificates from untrustworthy certificate issuers or a certificate authority that has been hijacked. Consequently, some phishing sites began to display HTTPS locks throughout the year.

When we examined the number of sites targeted by phishing attacks, we found that HTTPS hosting was used against every phishing target. The rate was highest for financial institutions; 87% of phishing attacks impersonating SunTrust Bank were HTTPS attacks, for BNP Paribas the figure was 66%, and for Wells Fargo 47%. Compare this with Google at just 17%. These numbers indicate that attackers can easily obtain SSL certificates and give users the “assurance” provided by the lock in the browser bar. The use of seemingly secure URLs using HTTPS, coupled with the growing popularity of phishing attacks, makes it even more important to use real-time anti-phishing technology to detect and stop phishing attacks before they can do harm.

THE “SWISS ARMY KNIFE” APPROACH TO ATTACKS

Bad actors continually learn from their mistakes and find new and cleverer ways to carry out cyberattacks. If one approach doesn’t work, another will, which is likely why we’ve seen an increase in what we call the “Swiss Army Knife” approach, when criminals employ multiple techniques to improve the likelihood that their attacks will succeed.

One of the best examples of this technique is Emotet. Emotet is a hugely successful polymorphic Trojan. CERT calls it “among the most costly and destructive malware affecting state, local, tribal, and territorial governments, and the private and public sectors.”^{ix} It has existed in various guises since 2014, having started out as a banking Trojan and evolved to use a variety of attack vectors. It infects a machine via spam email that uses branding the recipient will recognize and that contains a macro document. Recent campaigns have imitated PayPal receipts, shipping notifications, or past due invoices. Once the user opens the attached document and enables the macro, Emotet downloads and then spreads to other network-connected PCs by taking advantage of vulnerabilities and weak passwords. In particular, older or out-of-date systems or those that have not been patched provide a fertile ground for the botnet to spread. Emotet steals or brute-forces domain credentials, copying itself across the network and stealing email credentials in the process. It then sends itself to the next victim in a spam campaign.

93% of phishing domains in September and October 2018 offered an HTTPS site.^{viii}

SECURITY AWARENESS TRAINING PROVES ITS WORTH

An important ingredient in stopping the relentless onslaught of phishing attacks is security awareness training. The key to achieving dramatic results lies in the consistency and pace at which training is provided. Annual or semi-annual training won't achieve the desired results because phishers change their techniques and hooks from month to month, and security awareness training needs to keep pace with those changes and incorporate them into simulations and related training. In 2018, data from Webroot® Security Awareness Training, which combines training and phishing simulations to bring increased security and decreased risk, showed that organizations just starting to provide simulation campaigns saw a click-through rate of 29.32%. When combined with regular training, the rate can drop by 70%, a very significant decrease in the likelihood that an employee will click on a fraudulent link in a phishing email.

It's interesting to see that monthly training yields the best results; organizations that run 1-5 campaigns see a 28% click rate; after 6-10 campaigns the rate drops a bit to 24%, but the most impressive results are seen after 11 or more campaigns, where we saw an 8.5% click rate. Clearly platforms like Webroot Security Awareness Training can reduce the risk and costs of cyber threats through user education, presented at the time of greatest impact (i.e. immediately after clicking a link in a phishing simulation email).

After 12 months of training, end users are 70% less likely to fall for a phishing attempt.



Figure 10: Average number of clicks per number of campaigns

Emotet aims to increase the number of zombies in its spam botnet to gather credentials and act as a delivery vector for other malware. As discussed in the earlier polymorphic malware section, Emotet artifacts are typically found in arbitrary paths located off the %appdata% Local and Roaming directories.

Emotet exploded in capability and reach last year, hitting many corporate organizations around the world. Payloads were delivered at impressive speed, showing that bad actors have automated multiple steps in the campaign operations. They have also added layers to increase persistence and resilience. The Universal Plug and Play (UPnP) module lets Emotet turn victims' routers into proxy nodes for their command and control infrastructure. Criminals have also included additional email harvesting capabilities, which can now steal the contents of millions of Outlook emails on victims' machines.^x

A single malware campaign may combine phishing, spam, Trojans, botnets, cryptomining, vulnerability exploits, and more.

Several major malware campaigns use Emotet as a delivery vector. One is Trickbot, which is invariably seen with Emotet. Once Emotet has infected and spread, Trickbot can steal information for money. Trickbot has added Tor servers to its Level 1 command and control infrastructure, ensuring that the servers used to distribute the attack modules and web injections remain active for a longer period of time.^{xi} Trickbot also drops encrypting ransomware or other banking Trojans. Cryptominer payloads are another option, often dropped along with ransomware in case the user has backups, ensuring that there will be an alternate option for monetization.

Clearly, threat actors no longer rely on a single attack vector. They are making their malware more resilient, harder to detect, and longer-lasting.

MALICIOUS MOBILE APPS



With approximately 125 million Android™ smartphone users in the U.S. in 2017^{xii} and the preponderance of potentially harmful apps or PHAs (Google reports the probability of a smartphone user downloading a PHA is .02%^{xiii}), a large number of devices in the U.S. are infected. We saw a similar percentage in our analysis of hundreds of thousands of real-world Android users. In fact, our data showed an average of 0.22% - 0.62% of Android apps were malware and an average of 0.04% - 0.13% of Android apps were PUAs between September and December 2018.

Malicious apps operate in various ways, many of them designed to evade the defenses Google has put in place. An example is the TimpDoor campaign, which targeted U.S.-based phone numbers and devices and tried to turn them into network proxies. The campaign sent SMS messages to users informing them they had voice messages to review, but to hear them, users needed

to download a specific app. The app circumvented the security procedures and protections offered by the Google Play Store. When downloaded, the app enabled access to internal networks, bypassing firewalls and network monitors. At Webroot, we have seen many examples of this app on customer devices, showing that users continue to fall for these social engineering campaigns. Although Google strongly encourages users to download apps exclusively from trusted sources or the Google Play Store, there is an option to enable installation from “Unknown Sources”. Google found that users who have enabled installation from unknown sources see many more infections (an average of 0.65% in Q2, 0.86% in Q3, and 0.92% in Q4 in 2017) than those who only install from trusted sources and Google Play (0.08%, 0.12% and 0.09% for the respective quarters.)

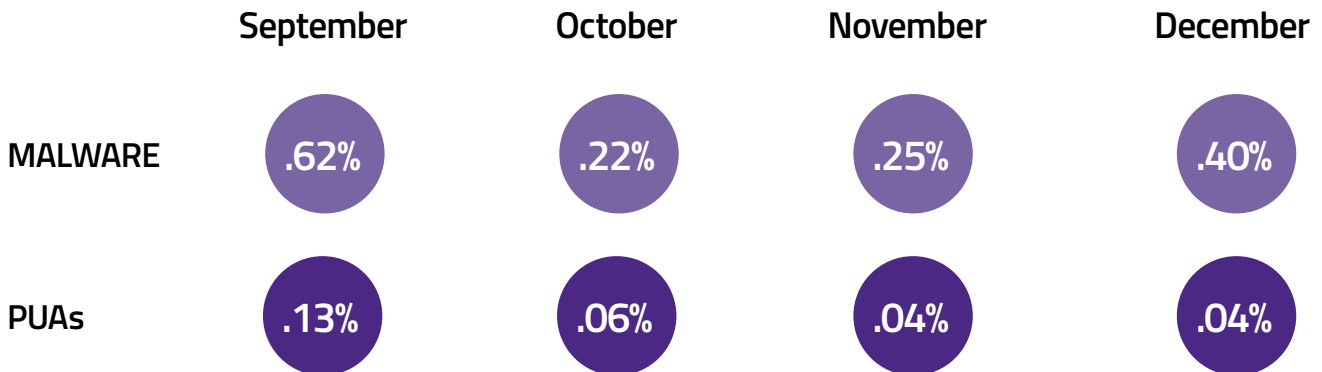


Figure 11: Percentage of apps that were malware or PUAs

SUMMARY AND PREDICTIONS FOR 2019



Threat actors continue to find new and innovative ways to combine attack methods for maximum results. Webroot data from 2018 shows that attackers can quickly pivot from one attack vector to another, with monetization a primary goal. The data we have collected and analyzed underscores the value of a multi-layered defense that stays up to date with the latest threats and approaches.

- » Polymorphic malware is here to stay: it continues to account for 93% of all malicious Windows executables.
- » While ransomware was less of a problem in 2018, it became more targeted. We expect major commodity ransomware to decline further in 2019; however, new ransomware families will emerge as malware authors turn to more targeted attacks, and companies will still fall victim to ransomware.
- » Cryptojacking and cryptomining surpassed ransomware in volume, providing quicker paths to profit despite dramatic fluctuations in the value of cryptocurrency. We expect that, in 2019, this trend will continue to dominate the landscape, with perhaps half of all attacks leveraging hardware in user devices to mine cryptocurrency.
- » High-risk IP addresses continue to be used and reused, with hundreds of thousands appearing on the black list at least two, three, or more times. While the majority are spam sites, a significant portion are botnets and scanners. Dynamically-updated IP address lists coupled with contextual analysis will continue to be the best way to deal with risky IP addresses.
- » We classified hundreds of millions of unique URLs in 2018, finding a high prevalence of phishing sites, spam sites, and botnets. We found that 40% of malicious URLs were on good domains, and users can't easily tell whether the destination is benign or malicious. We expect this trend to continue in 2019. To protect users, solutions need URL-level visibility or better domain-level metrics that accurately represent the dangers.
- » We foresee an increase in router and IoT targets in 2019. We've already seen a major instance of this, as more than 400,000 MikroTik routers around the world were hijacked to conduct Monero mining.^{xiv} Home users are not immune; their routers serve as the hub for networks and smart home devices (IoT), yet most users can't log into their Linux-based routers to see what they are doing. Meanwhile a hacker can learn everything about a user's environment, can redirect URLs, carry out man-in-the-middle attacks, and even inject cryptojacking scripts.



I anticipate Monero will continue becoming the cryptocurrency of choice for cybercriminals because of its private ledger and fungibility. The SamSam criminals were indicted because the FBI was able to track all ransom addresses to the exchange. Bitcoin is not anonymous and criminals will adapt to this fact.

TYLER MOFFITT | SECURITY ANALYST



» The sophistication, agility and innovation of threats—malware, malicious IPs and URLs, cryptojacking and cryptomining, malicious mobile apps, targeted phishing, ransomware attacks, and others that rely on social engineering to exploit human nature—all point to the fact that today’s defenses must be constantly updated. Just as attackers use the Swiss Army Knife approach, organizations must employ a multi-layered security strategy. Automated, real-time tools based on continuously-updated threat intelligence, contextual analysis and advanced endpoint and network protection are a must, as is timely security awareness training that is constantly updated to help users avoid the latest threats. This combination of tools and training can help any organization materially reduce its exposure to unacceptable risk.

ABOUT THE DATA

The statistics presented in this annual threat report are derived from metrics automatically captured and analyzed by the Webroot® Platform, our advanced, cloud-based machine learning architecture. This system provides proactive protection for users and networks against both known and zero-day, never-before-seen and advanced persistent threats. Threat intelligence produced by the platform is used by Webroot® endpoint security products and by technology partners through Webroot BrightCloud® Threat Intelligence Services. Our threat intelligence is based on visibility of the entire IPv4 and in-

use IPv6 space, billions of URLs, tens of millions of new and updated mobile apps, and all Webroot-protected endpoints worldwide. Advanced machine learning techniques, real-time scoring with confidence levels, and continuous updates enable Webroot threat intelligence to be highly effective at identifying and stopping even the most sophisticated threats. Webroot takes a unique approach to machine learning, based on massive data processing capacity, a proprietary implementation of the most advanced technology available, and a powerful contextual analysis engine. Contextualization is a “guilt by association” model that links internet objects. Capturing an extensive range of characteristics for each internet object observed (up to 10 million characteristics per object) enables Webroot to determine if the object poses a threat at the precise time of analysis. Our patented approach maps attack and threat behavior across vectors, analyzing the relationships among URLs, IPs, files and mobile apps. For example, if a user runs a mobile app that tries to access the contact list and transfer it to an IP address, the malicious behavior of the app would impact the reputation score of the IP address. This ability to correlate current associations among objects with history on how millions of objects have behaved over time is what makes Webroot threat intelligence predictive in nature.

ⁱ www.appsteem.com

ⁱⁱ US hospital pays \$55,000 to hackers after ransomware attack. ZDNet, Jan 2018. Retrieved from www.zdnet.com/article/us-hospital-pays-55000-to-ransomware-operators

ⁱⁱⁱ Wanted by the FBI: SamSam subjects. FBI, November 2018. Retrieved from www.justice.gov/opa/press-release/file/1114746/download

^{iv} Town, Sam. Physical Bitcoin Attacks and Burglaries on the Rise. CryptoSlate, July 2018. Retrieved from cryptoslate.com/physical-bitcoin-attacks-and-burglaries-on-the-rise

^v \$731 Million Stolen from Crypto Exchanges in 2018: Can Hacks be Prevented? Bitcoin Exchange, July 2018. Retrieved from www.ccn.com/731-million-stolen-from-crypto-exchanges-in-2018-can-hacks-be-prevented

^{vi} S. Pastrana, G. Suarez-Tangil. A First Look at the Crypto-Mining Malware Ecosystem: A Decade of Unrestricted Wealth. Cornell University, January 2019. Retrieved from arxiv.org/abs/1901.00846

^{vii} A. Newmann, J. Barnickel, U. Meyer. Security and Privacy Implications of URL Shortening Services. RWTH Aachen University, 2011. Retrieved from pdfs.semanticscholar.org/e9a0/8f1bc561310db382bdacbb397b4ca005b9bb.pdf

^{viii} 2018 Phishing and Fraud Report: Attacks Peak During the Holidays. F5 Labs, Nov. 2018. Retrieved from [2018 Phishing and Fraud Report: Attacks Peak During the Holidays](http://2018-Phishing-and-Fraud-Report-Attacks-Peak-During-the-Holidays)

^{ix} Alert (TA18-201A). US-CERT, United States Computer Emergency Readiness Team. July 2018. Retrieved from www.us-cert.gov/ncas/alerts/TA18-201A

^x Emotet Malware gang is mass harvesting millions of emails in mysterious campaign. ZDNet, Oct 2018. Retrieved from www.zdnet.com/article/emotet-malware-gang-is-mass-harvesting-millions-of-emails-in-mysterious-campaign

^{xi} Advisory: Trickbot Banking Trojan. National Cyber Security Centre, Sept 2018. Retrieved from www.ncsc.gov.uk/alerts/trickbot-banking-trojan

^{xii} Insights into the 2.3 Billion Android Smartphones in Use Around the World. New Zoo, Jan 2018. Retrieved from newzoo.com/insights/articles/insights-into-the-2-3-billion-android-smartphones-in-use-around-the-world/

^{xiii} Google Android Security report 2017: We read it so you don't have to. ZDnet, Mar. 2018. Retrieved from www.zdnet.com/article/google-android-security-report-2017-we-read-it-so-you-dont-have-to-and-here-are-the-takeaways

^{xiv} 415,000 routers infected by cryptomining malware – prime target MikroTik. ETHNews, Dec. 2018. Retrieved from www.ethnews.com/researchers-claim-400-000-mikrotik-routers-infected-with-mining-malware

About Webroot

Webroot was the first to harness the cloud and artificial intelligence to protect businesses and individuals against cyber threats. We provide the number one security solution for managed service providers and small businesses, who rely on Webroot for endpoint protection, network protection, and security awareness training. Webroot BrightCloud® Threat Intelligence Services are used by market leading companies like Cisco, F5 Networks, Citrix, Aruba, Palo Alto Networks, A10 Networks, and more. Leveraging the power of machine learning to protect millions of businesses and individuals, Webroot secures the connected world. Headquartered in Colorado, Webroot operates globally across North America, Europe, and Asia. Discover Smarter Cybersecurity® solutions at webroot.com.

385 Interlocken Crescent Suite 800 Broomfield, Colorado 800.870.8102 webroot.com