


proofpoint.

A high-angle photograph of three people sitting on wide, light-colored stone steps. The person on the left is a man with dark hair, wearing a brown jacket and dark pants, looking at a tablet. The person in the middle is a man with a balding head, wearing a tan sweater and blue jeans, looking at a smartphone. The person on the right is a woman with dark curly hair, wearing a bright yellow jacket, looking at a smartphone. The steps are made of large, rectangular stone tiles. The background is a continuation of the steps, leading up and away from the camera.

HUMAN FACTOR REPORT

2019

THE HUMAN FACTOR 2019

Each year, cyber criminals continue to refine their use of social engineering, relying on human interaction rather than automated exploits to install malware, initiate fraudulent transactions, steal data and engage in other malicious activities. Less than 1% of the attacks we observed made use of system vulnerabilities. The rest exploited “the human factor”: the instincts of curiosity and trust that lead well-intentioned people to click, download, install, open and send money or data.

Instead of attacking computer systems and infrastructure, threat actors focused on people, their roles within an organisation, the data to which they had access, and their likelihood to “click here.” Whether attacking at a massive scale in large indiscriminate campaigns, going after specific industries or geographies with more targeted campaigns, or seeking out a single person within an organisation, attackers and their sponsors consistently found human beings to be the most effective vectors to infiltrate organisations and facilitate fraud and theft.

While ransomware was the biggest threat of 2017, the last 18 months have seen a marked shift towards information-stealing malware, with social engineering becoming ever more pervasive and effective at preying on people. Whether sending impostor messages that appear to come from a trusted colleague or installing increasingly robust malware that can silently profile individuals and steal data and credentials to make future attacks more effective, threat actors are following the money. While cryptocurrency volatility and a growing ability to detect and mitigate ransomware may have driven this shift initially, the information provided by victims via malware and phishing attacks is fuelling revenue streams and facilitating future attacks.

Regardless of the means of attack—email, cloud applications, the web, social media, or other vectors—threat actors repeatedly demonstrated the effectiveness of the social engineering tactics that convinced victims to click malicious links, download unsafe files, install malware, transfer funds and disclose sensitive information at scale. Whether financially motivated or state-sponsored, attackers all had one thing in common: an understanding of and a willingness to take advantage of the human factor.

TABLE OF CONTENTS

Key Findings	4
Attackers target people and not necessarily who you might expect	4
Social engineering reaches critical mass	4
Vectors multiply as attackers refine techniques and go after a range of platforms	5
Introduction	5
By the Numbers	6
Top 20 phishing lures	6
Click rates for top 20 most-clicked phishing lures	6
Click times by region	7
Average number of impostor attacks per targeted company by industry	8
Top impostor email subject lines	9
Impostor email subject lines for the top five most targeted industries	9
Impostor message volume by day of week	10
Change in impostor attack vector by quarter	11
Top 10 industries targeted by malware campaigns	12
Relative distribution of malware strains for the top five targeted industries, 2018	12
Finance	13
Manufacturing	13
Technology	13
Healthcare	13
Retail	13
Relative distribution of attack techniques for the top five targeted industries, 2018	14
Finance	14
Manufacturing	14
Technology	14
Healthcare	14
Retail	14
People-centred attack metrics	15
Attack Index by industry	15
“Very Attacked People” by industry	16
VAP sources	17
Tools and Techniques	18
Social engineering basics	18
Domain abuse: Brand theft and look-alikes	20
Good for bad: Leveraging legitimate infrastructure	22
Impostor attacks	22
Conclusion	23
Recommendations	24

KEY FINDINGS

Email remains the top attack vector. Threats range from malicious spam that clogs inboxes and wastes resources to impostor attacks that can cost organisations and people millions of pounds. Threat actors attack cloud applications, leverage increasingly robust multi-purpose malware and seek out new ways to steal both money and data directly.

Here are the key findings from Proofpoint research over the 18 months of 2018 and the first half of 2019. The results, based on data collected across our global customer base and analysis of billions of messages per day and hundreds of millions of domains, highlight the ways in which actors are increasingly exploiting “the human factor.”

Attackers target people and not necessarily who you might expect

The modern threat landscape is increasingly “people-centric.” Attacks focus on people and identities rather than infrastructure, making it more important than ever to identify which users in an organisation represent the greatest sources of risk.

- “Very Attacked People” (VAPs) represent significant areas of risk for organisations. They tend to be either easily discovered identities or targets of opportunity like shared public accounts. Of the identified VAPs, 36% of the associated identities could be found online via corporate websites, social media, publications and more.
- VAPs are not necessarily high-profile individuals either (VIPs such as C-level executives). Only 7% of executive emails could be found online.
- For the VIPs who are also VAPs, almost 23% of their email identities could be discovered simply by a Google search.
- Education, finance, and advertising/marketing were the industries with the highest average Attack Index, an aggregated measure of attack severity and risk.

Social engineering reaches critical mass

Attackers are increasingly focused on obtaining credentials to feed further attacks and are improving the social engineering techniques with which they obtain them. Similarly, malware distribution is far more focused on establishing a silent foothold in organisations to commit fraud and steal data and credentials rather than simply smash-and-grab via ransomware attacks.

- Generic email harvesting accounted for almost 25% of all phishing schemes in 2018. In 2019, Microsoft Office 365 phishing has been the top scheme, but the focus remains credential harvesting.
- The most effective phishing lures in 2018 were dominated by “Brain Food,” a diet and brain enhancement affiliate scam that harvested credit cards. However, 2019 saw a shift in terms of effectiveness towards cloud storage, DocuSign and Microsoft cloud service phishing.
- Impostor attacks include schemes like business email compromise (BEC) and also include increasingly mainstream identity deception techniques used in a variety of scenarios supporting social engineering and more effective people-centred campaigns. 2018 saw impostor attacks at their highest levels in the engineering, automotive and education industries. This likely reflects easily exploited supply chain complexities in the first two and high-value targets and user vulnerabilities—especially among student populations—in the latter.
- Over 99% of emails distributing malware required human intervention—following links, opening documents, accepting security warnings and other behaviours—for them to be effective.

Vectors multiply as attackers refine techniques and leverage a range of platforms

The ways in which attackers select and target potential victims multiply as attackers refine techniques and leverage a range of platforms:

Themes vary widely by both actor and intended target. Food, shelter, love and money are perennial favourites, feeding everything from threats of lawsuits over food poisoning in attacks on restaurants to rampant sextortion schemes targeting individuals.

- BEC tactics—building rapport with attacked individuals, multiple points of contact and creating a sense of urgency, among others—began appearing more frequently in attacks involving commodity malware.
- Domain fraud and abuse ramped up even more, with attackers leveraging a range of techniques from look-alike domains to legitimate secure certificates to make malicious websites appear trustworthy.

INTRODUCTION

Malware-free attacks like business email compromise (BEC) and credential phishing continue to rapidly gain momentum as threat actors consistently attack individuals and business processes rather than specific systems and software. Similarly, attacks on Software-as-a-Service (SaaS) accounts and platforms create new risks for businesses that are increasingly reliant on the cloud. At the same time, malware-based attacks have shifted almost entirely to payloads that support long-term credential theft, information gathering, and the ability to flexibly load new malware in the future—all while evading detection. This pendulum swing toward persistent, non-destructive infection and information theft allows attackers to collect more and more data about us, all of which can be turned around in a feedback loop for a range of highly targeted attacks.

This year, we are uniquely positioned to examine four key components of the attacks businesses and individuals face every day:

- Who within an organisation is being attacked?
- How are phishing and impostor attacks (including BEC) evolving?
- What malware is being used to attack individuals and organisations? And how is it being targeted and used?
- What kinds of attacks are increasingly targeting the SaaS platforms on which people and businesses rely? And how does the trend towards pervasive credential phishing (both socially engineered and malware-driven) feed into this?

BY THE NUMBERS

Top 20 phishing lures

Phishing lures leverage a range of brands and are set to resemble login portals from banks, online retailers, webmail and more. While some are phishing for credentials from specific services, it appears that many are simply looking for the email logins used with various services to inform credential-stuffing attacks. By far, the most prevalent phishing campaigns in 2018 sought a variety of email login credentials, often offering to allow users to log in to fake services with any number of email accounts. This type of generic email harvesting accounted for almost 25% of all phishing schemes.




Figure 1: 2018 relative phishing campaign volumes¹

CHALBHAHAI

Chalbhai phishing refers to a range of campaigns associated with templates created and sold by a group collectively referred to as Chalbhai based on artefacts in the templates.

In the first half of 2019, the focus has remained on email credential phishing, although Microsoft Office 365 account phishing took the top spot from more generic efforts. So-called “Chalbhai phishing” took the third spot. **CHALBHAHAI** phishing targets credentials for a number of top US and international banks, telecommunications companies and Microsoft services (among others), using a range of templates attributed to a single group but leveraged by multiple actors.

Click rates for top 20 most-clicked phishing lures

As we noted in previous years, the most-clicked lures are not necessarily the same as the most common phishing lures. In 2018, Brain Food phishing was by far the most clicked lure, with click rates over 1.6 **CLICKS PER MESSAGE**. This indicates that the messages themselves are shared and clicked through multiple times on many occasions. Brain Food refers to a botnet that distributes diet and mental enhancement spam and routinely sends users to credit card harvesting pages among others. Blackboard phishing, which collects credentials for the popular school management system, WeTransfer (a file-sharing service), and Zoominfo (a business contact database) all had click rates above 0.6 clicks per message.

CLICK RATE

The number of times a single message in a campaign entices a recipient to click a link. Click rates greater than 1 suggest that a message may have been forwarded frequently or the URL in the message was revisited multiple times on average across the campaign.

IN THE FIRST HALF OF 2019, WE HAVE OBSERVED A SHIFT TOWARDS CLOUD STORAGE LURES IN TERMS OF EFFECTIVENESS.




Figure 2: 2018 average click rates for phishing messages from the top 20 most-clicked lures

Click rates were nearly identical across industries, between 20% and 22% for all lures in aggregate. However, in the first half of 2019, we have observed a shift towards cloud storage lures in terms of effectiveness—DocuSign, and Microsoft cloud service phishing in particular—with Brain Food activity dropping off and educational phishing often loaded towards the third quarter as school begins.

Click times by region

Click times have traditionally shown significant regional differences, reflecting differences in work culture and email habits among major global regions. Figure 3 shows the percentage of phishing links clicked by time of day, with Asia-Pacific and North American organisations far more likely to read and click early in the day and Middle Eastern and European users more likely to click midday and after lunch. This is important for defenders in organisations looking to enhance training, mitigation and end-user outreach around phishing. Click times were averaged across five quarters, as they showed little quarter-to-quarter variation. These habits appear relatively fixed.




Figure 3: Relative click times by region, adjusted for local times, Q1 2018-Q1 2019

IMPOSTOR ATTACK
Impostor attacks rely on a range of identity deception techniques to trick users into completing an action. These techniques can include domain spoofing, the use of look-alike domains and more.

Average number of impostor attacks per targeted company by industry

2018 saw **IMPOSTOR ATTACKS** at their highest levels in the engineering, automotive and education industries, likely reflecting easily exploited supply chain complexities in the first two and high-value targets and user vulnerabilities, especially among student populations, in the latter.




Figure 4: Impostor attacks by industry relative to the weighted global average impostor attacks per targeted organisation

To date in 2019, we have observed a shift in the most highly targeted industries, with financial services, manufacturing, education, healthcare and retail rounding out the top five. Regardless of industry, however, there continues to be no correlation between the size of targeted companies and the number of impostor attacks they receive. While larger organisations may be attractive for their deep pockets, smaller companies may be more vulnerable due to relative lack of controls and awareness, both of which create lucrative potential outcomes for threat actors.

Top impostor email subject lines

Common subject lines in impostor emails continue to shift away from those related to a “Request” and towards “Payment” scams. Similarly, “Urgent” subjects are showing an upward trend while W2-related attacks maintained their expected seasonal upticks at the end of 2018 and beginning of 2019, corresponding to tax processing seasons.




Figure 5: Most common subject line categories in impostor emails indexed by quarter, Q1 2018 to Q1 2019

Impostor email subject lines for the top five most targeted industries

Subject lines associated with impostor emails not only vary seasonally but also by industry. Looking more closely at the top five targeted industries, we see that education receives a disproportionate number of impostor emails related to “Request” and “Greeting” while attacks on engineering firms, for example, favour “Urgent” and “Request” messages. In any of these cases, the subject lines associated with impostor attacks are often tailored to the receiving industry in ways that more traditional **AFFILIATE** or malware-bearing spam is not.




Figure 6: Top impostor email subject categories broken out by the five most targeted industries

Impostor message volume by day of week

AFFILIATE SPAM

A type of spam that entices users to visit pages for pharmaceuticals, work-from-home schemes, books, dating sites, or any number of services that pay affiliates in multi-level marketing schemes for clicks through to their sites.

Impostor message delivery closely mirrors trends observed in previous years, with less than 5% of overall messages delivered on weekends and the largest portion—over 30%—delivered on Mondays. Threat actors capitalise on Monday morning backlogs and social jetlag to more easily fool people with impostor tactics and other social engineering elements. Delivery rates steadily drop off through the week, and delivery rates are nominal when few employees are in the office.




Figure 7: Relative volume of messages associated with impostor attacks by day of the week

By way of comparison, overall malicious message volumes sampled in the second quarter of 2019 were distributed more evenly over the first three days of the week and were also present in significant volumes in campaigns that began on Sundays, suggesting that malware actors in general are less concerned about timing than actors focusing on impostor techniques and the inherent social engineering approaches.




Figure 8: Relative volume of messages associated with all malware attacks by day of the week, Q2 2019

SPOOFING

Forgery techniques that alter the appearance of email headers such that they seem to be from someone else.

Change in impostor attack vector by quarter

In general, since the beginning of 2018, we have observed an ongoing and steady increase in the frequency with which multiple **SPOOFED** identities are used to target many individuals in organisations. For example, threat actors might spoof the identities of several executives or senior managers, sending a malicious document to a range of employees asking them to read the document. Alternatively, in a more traditional BEC-style attack, several spoofed identities might be used to ask all staff in the Human Resources department to forward W2 information for employees.

While so-called one-to-one and one-to-many attacks were more common when impostor attacks first began to emerge—especially in the context of business email compromise—threat actors appear to be finding success in attacks using more than five identities against more than five individuals in targeted organisations. One point to note, however, is the relative increase, albeit at a smaller scale, in attacks over the last two quarters in which a single spoofed identity was used to target a single individual.




Figure 9: Relative volume of impostor attacks by type and quarter, Q1 2018 to Q1 2019

Top 10 industries targeted by malware campaigns

As with phishing, the top industries targeted by malware actors vary from year to year. However, financial services, manufacturing, technology, healthcare and retail frequently top the list, as they did in 2018.




Figure 10: Relative total message volumes for malware campaigns in the top 10 most-targeted industries

Relative distribution of malware strains for the top five targeted industries, 2018

Banking Trojans, information stealers, downloaders and botnets regularly appeared in malware campaigns in 2018 as threat actors continue to favour longer-term information gathering and credential theft over the more damaging but short-term gains associated with ransomware that was endemic in 2016 and 2017. While data presented are from 2018, we have observed similar trends in the first half of 2019.

Note that while the general distribution of malware in aggregate favours banking Trojans, **REMOTE ACCESS TROJANS (RATS)** and downloaders, individual industries experienced different proportions of malware families. In some cases, specific malware strains were only common in particular industries, based on the modus operandi of the actors targeting the verticals and specific vulnerabilities frequently associated with those industries. For example, the financial services sector tended to see higher volumes of **FLAWEDAMMY** and **SERVHELPER** as the actor we track as TA505 turned their attention to finance. Technology companies, on the other hand, saw higher proportions of downloaders and banking Trojans, potentially for both financial gain and to load software for stealing intellectual property or launching further attacks. Again, the mail exchanger records (MX records) are noteworthy, as fraudulent domains appear to be coming in line with all domains, although it is not clear why threat actors are now more likely to create MX records for their domains.

REMOTE ACCESS TROJANS (RATS)

Also known as a Remote Administration Tools, RATs are robust malware that can control most elements of an infected host and provide a range of data to threat actors.

FLAWEDAMMY

A remote access Trojan based on leaked source code of the AmmyyAdmin remote administration software.

SERVHELPER

A backdoor with two variants: One that primarily acts as a RAT and another frequently deployed as a downloader.



Figure 11: Relative malicious message volumes by payload for the top five most-targeted industries

AZORULT

A robust information stealer originally discovered in 2016.

URSNIF

A banking Trojan frequently deployed in global and regional campaigns by a range of affiliates.

We have continued to see strong targeting of these industries in 2019. Looking at the breakdowns for the top five gives a sense of how threat actors adapt techniques for various industries in some cases (such as the use of **AZORULT** in financial services). In other cases, threat actors still send broad-based campaigns (such as **URSNIF**) that affect organisations across a range of industries.

Relative distribution of attack techniques for the top five targeted industries, 2018

Regardless of the malware being distributed, 99% of the malware we observe requires at least some degree of human interaction to infect user devices. With exploit kits continuing to operate at a tiny fraction of their 2016 peak and many software vulnerabilities quickly addressed by more aggressive vendor patching, campaigns throughout 2018 and the first half of 2019 relied on users to click links, open documents, enable macros, bypass security alerts or unzip malicious executables. In the face of increasingly effective social engineering, people continue to do just that.

Figure 12 shows the most common attack techniques observed in the industries with the highest malware volumes in 2018. Microsoft Office Visual Basic for Applications (VBA) macros are the common thread. Malicious macro-laden documents are the most common vector we observe in aggregate when hosted and distributed with a link or attached to email. However, even the specific vulnerabilities noted below, like CVE-2017-11882, require humans to open the malicious document that leverages the exploit.



Figure 12: Relative malicious message volumes by leveraged technique for the top five most-targeted industries

People-centred attack metrics

Attacker strategies go well beyond social engineering and convincing people to click. Threat actors are increasingly focused on getting the right people to click—those with sufficient access and privileges to best establish a beachhead on a network or with those most likely to move funds in a wire transfer scheme, for instance. Attackers also focus on “targets of opportunity,” often going after shared accounts that are difficult to secure or accounts with large public and social media footprints. The analyses below quantify the scale and targeting of these so-called “people-centred attacks.”

Attack Index Defined

The Attack Index is an aggregate measure of cybersecurity risk for individuals in an organisation. Measuring it by person allows businesses to allocate security resources to individuals and departments with the greatest risk of an effective or damaging attack. Looking at summary measures of the index across industries provides a better understanding of the threat landscape, threat actor behaviour and the challenges faced by specific industries over time.

The Attack Index is based on three components:

- Actor type
- Targeting type
- Threat type

Actor type considers the attacker's level of sophistication. For example, an advanced persistent threat (APT) state actor will be given a higher score than a small-scale, financially motivated crimeware actor.

Targeting type speaks to the degree of targeting involved with the threat. Did the threat hit only a small number of global users? Was it focused on a particular user, company, industry or geography? Or was it a spray-and-pray campaign seen by half the globe? The former will receive a higher score than the latter.

Threat type addresses the type of malware involved in the attack. This considers how dangerous the threat is and how much effort went into the threat. In this case, a RAT or stealer is going to have a higher score than a generic consumer credential phish.

Attack Index by industry

As noted, attackers are increasingly focused on attacking the “right people” in an organisation rather than attacking every user and seeing which attacks are successful. These “right people,” whether targets of opportunity or identified users with sufficient access and privilege, generally make up groups of VAPs in an organisation. Looking at average Attack Index and the average number of VAPs across industries provides a view of both the overall severity of the threats against organisations in a given industry and the number of accounts within industries that threat actors are able to identify for targeted attacks.

Figure 13 shows that education, for example, is frequently targeted with attacks of the highest severity and has one of the highest average number of VAPs across industries. Financial services, on the other hand, has a relatively high average Attack Index, but appears to do a better job concealing individual identities and accounts from attackers, creating fewer VAPs to attack.




Figure 13: Average Attack Index associated with major industry groups versus the average number of VAPs per organisation in the industry

“Very Attacked People” by industry

Figure 14 provides a longitudinal view of the changing number of VAPs per month from September 2018 through February 2019. Seasonal dips around the holidays and the end of the year correspond to fewer targeted attacks overall during this period, while most industries showed an upward trend in the average number of VAPs targeted monthly coming into 2019. Education, heavy industry and healthcare consistently topped the list with the highest numbers of VAPs calculated monthly.




Figure 14: Average number of VAPs per organisation by industry

VAP sources

Interestingly, VIPs like C-level executives and board members are often not VAPs. Rather, email addresses for VAPs tend to be more easily identified online than those for VIPs, making it simpler for attackers to discover their contact information and role and target them with high-severity threats. Of the VAPs we examined in a sample across industries, 36% of them could be found online. Note that some VAPs may have been found in more than one place.




Figure 15: Common sources of VAP identities, 2018




Figure 16: Common sources of VIP identities, 2018

For executives in our sample, only 7% of their email addresses could be discovered online.

However, for the VIPs who are also VAPs, almost 23% of their email identities could be discovered simply by a Google search. The intersection of VAP and VIP represents an area of particular risk for organisations.

THE INTERSECTION OF VAP AND VIP REPRESENTS AN AREA OF PARTICULAR RISK FOR ORGANISATIONS.

SOCIAL ENGINEERING APPROACHES RANGE FROM SIMPLE LURES DESIGNED TO SPARK SUFFICIENT CURIOSITY FOR VICTIMS TO OPEN A MALICIOUS DOCUMENT ATTACHMENT TO MUCH MORE ELABORATE SCHEMES.

GRIFFON

A JavaScript backdoor frequently leveraged by the Carbanak/FIN7 group in their financially motivated campaigns.

TOOLS AND TECHNIQUES

Exploiting the human factor is the bread and butter of modern threat actors. Convincing people to take action—whether revealing credentials and sensitive information through phishing attacks or installing malware by following malicious links or opening macro-laden documents—relies on a range of tools and techniques. Each of these techniques builds rapport with individuals, creates believable situations, establishes credibility and creates a sense of urgency. Several categories of these techniques are outlined below.

Social engineering basics

Social engineering is at the core of the majority of attacks we observe. Approaches range from simple lures designed to spark sufficient curiosity for victims to open a malicious document attachment—for example, a fake invoice sent to an accounts payable team or a résumé emailed to Human Resources staff—to much more elaborate schemes. These might include threat of a lawsuit over a fabricated incident or threat through exposure of potentially damaging online habits.

While themes vary widely by both actor and intended target, food, shelter, love and money are perennial favourites. Carbanak Group, for example, is a sophisticated financially motivated actor who uses carefully crafted lures and professional-looking document attachments to distribute multiple strains of malware. For example, a Carbanak campaign in late 2018 featured a document attached to an email in which the sender claimed to have been double-charged and demanded resolution. The document used an increasingly common tactic among threat actors: stolen security vendor branding with claims that the document was protected by the vendor's technology. Embedded instructions to “decrypt” the document, however, were actually the steps to enable macros and allow the installation of malware. In the Carbanak example, these macros installed the **GRIFFON** backdoor. We have observed similar campaigns abusing a range of security company brands for social engineering.

This type of social engineering in a convincing lure and a realistic malicious document are typical of Carbanak and often carries into targeting as well. For example, Carbanak frequently targets restaurant chains with lures related to purported cases of food poisoning at a restaurant.

Real estate lures also provide noteworthy examples of social engineering. Because real estate transactions often involve contact with multiple parties, high degrees of urgency, the exchange of personal information and digital signatures, they are frequent targets for criminals engaged in both phishing and malware attacks. Threat actors routinely abuse the DocuSign brand, a trusted source for electronic signatures, at all steps of a real estate transaction. In addition to the frequent abuse of trusted brands like DocuSign and various realty and bank portals, the stress and often unknown elements associated with buying a home or applying for a rental create powerful opportunities for threat actors to leverage the human factor.

Love and sex, on the other hand, prey on the lonely and general concerns about privacy. Dating scams, affiliate spam related to various products, and so-called “sextortion” scams are rampant in email. In sextortion schemes, threat actors send messages claiming to have evidence of the victim’s potentially damaging online activities, often backing up their claims with passwords associated with the recipient’s email account or publicly available personal information. While the passwords themselves may be guesses or old passwords gleaned from a data breach, the emails are designed to create panic and convince users to quickly pay the sender not to reveal browsing history, compromising webcam photos and more. In most cases, these are straight blackmail schemes, but we have also detected malware attacks delivered with the scams as well.

Regardless of the particular theme, some degree of social engineering appears in most campaigns. Given that over 99% of the threats we observe require human interaction to execute—enabling a macro, opening a file, following a link, or opening a document—the element of social engineering is key to successful attacks. More importantly, this technique is particularly effective, taking advantage of human vulnerabilities when software vulnerabilities are increasingly rare.

Actor Spotlight: TA557’s Fake Jobs

An actor we track as TA557 uses multiple points of contact to complete their social engineering scheme and better establish a relationship with the victim. TA557 sends a LinkedIn invitation to the victim using a legitimate account and then follows up with a personalised email without any malicious content. It is not until a subsequent email that the actor sends a malware-bearing message.

This type of attack, with multiple touch points and deep social engineering, has previously been associated with BEC, a scenario where actors convince victims with access to corporate funds or sensitive information to initiate fraudulent wire transfers or send sensitive personal information to a threat actor posing as a business leader in a position of authority. In this case, however, TA557 uses the interactions to more effectively trick victims into installing malware. We are increasingly observing this cross-pollination of techniques and more effective targeting by malware actors who are capitalising on the human factor rather than the scale of their attacks.

Domain abuse: Brand theft and look-alikes

Believable domains and web presence tangibly support social engineering efforts. Whether malicious and fraudulent websites use stolen branding to create legitimate-looking landing pages or threat actors register domains that resemble those of real brands, fraudulent domains are a critical piece of the cyber criminal's toolkit.

While the frequency with which particular brands are abused changes from month to month, the first half of 2019 bore perennial favorites. Threat actors stole branding and created lures and landing pages most often for

- Major US and international banks
- Amazon and other major retailers
- Cellular providers
- Shipping carriers
- Document signing and electronic faxing services

Hundreds of brands, however, frequently appeared in malicious emails.

Although many threats simply stole visual branding without regard to hosting domains, many others layered domain fraud techniques into their scams. Look-alike domains are increasingly sophisticated, sometimes making use of Unicode characters that render indistinguishably from the ASCII characters for which they are substituted. In other cases, threat actors rely on more traditional approaches, substituting the number three for the letter “E”, the number one for the letter “l”, and so on. Regardless of the particular substitutions, at a glance, the domains are close enough to the originals to fool even savvy but hurried victims.

Beyond strict look-alikes, threat actors also register seemingly legitimate variations of brand-owned domains. For example, acmeanvilssupport[.]com would appear to be associated with acmeanvils[.]com but provides opportunities for social media support account fraud (also known as angler phishing), impostor email attacks and more.

Notably, these domains feature secure certificates at much higher rates than domains across the web, creating a false sense of security and privacy for potential victims. Recent Proofpoint research demonstrated that fraudulent domains implemented **SECURE SOCKETS LAYER (SSL)** certificates at three times the rate of their legitimate counterparts. At the same time, threat actors tended to register .com domains for these schemes, further adding to the sense of trust and familiarity that cannot be achieved with **TOP-LEVEL DOMAINS (TLDs)**.

SECURE SOCKETS LAYER (SSL)

Encryption technology used to ensure that communication between a web browser and a web server cannot be intercepted by third parties.

TLD

Top-level domains refer to the .com, .net, .biz, and other broad domain types administered by the international body ICANN and available to register based on specific requirements and organisation types.




Figure 17: Top TLDs appearing in fraudulent domains 2018. Note that “.xn--p1ai” is the ASCII equivalent of the Cyrillic Unicode text representing the .ru TLD

Good for bad: Leveraging legitimate infrastructure

As we increasingly move infrastructure to SaaS platforms and rely on a variety of file-sharing, collaboration and communications tools, threat actors take advantage of their familiarity and frequent whitelisting to distribute malware, host phishing templates and more. Frequently abused platforms include:

- Document collaboration services like Google Drive and Microsoft Office 365
- File-sharing services like Box and Dropbox
- Mass mailing services like MailChimp and SendGrid
- Payment services that allow outbound mailing of invoices
- Social media platforms

Recently, we detected sophisticated phishing templates **hosted on GitHub**, a ubiquitous platform for code development that includes free hosting for projects, while campaigns hosting malware on Google Drive, Microsoft SharePoint and other similar services are almost daily occurrences. Not only can these attacks bypass traditional defences because legitimate uses prevent organisations from blacklisting them, but they also leverage the human factor quite effectively. Because we are conditioned to open links received in notification emails from these services, users often do not consider that the links may lead directly to malware or to credential phishing.

In fact, although the human factor is focused on attacks against people rather than infrastructure, SaaS infrastructure is the exception that makes the rule. If attackers can infiltrate SaaS platforms, they can launch a range of secondary attacks like those described above that are hard to detect algorithmically and even harder to spot by users. In addition to using the platforms to send spam or host malicious content, they can conduct internal phishing that require sophisticated detection mechanisms to separate from legitimate emails from trusted colleagues. The widespread availability of credential dumps has also provided threat actors with deep data mining capabilities that inform credential-stuffing attacks and other intelligent brute-force attacks on SaaS platforms.

Equally important, people often reuse passwords—another element of the human factor—making even badly dated credential dumps useful sources of information for attackers looking to compromise SaaS accounts and platforms. Recent Proofpoint research discovered that 45% of organisations have at least one compromised account and 6% have at least one VIP account compromised, making effective internal phishing and BEC relatively easy for attackers.

Impostor attacks

Impostor attacks leverage a range of techniques to convince victims that they are actually communicating with a trusted entity. These include display-name spoofing, domain spoofing, and look-alike domains and may lead to wire transfer fraud, angler phishing, malware attacks and more. The unifying attribute, however, is “identity deception,” where threat actors pretend to be a CEO, a colleague, a business partner, or even support staff for a given brand interacting with customers. These are differentiated from more common attacks that simply use throwaway attacker-owned addresses and domains to launch and host malware and phishing attacks.

While impostor attacks were initially most closely associated with BEC, they are proving highly effective tools for a range of payloads and attacks. If employees receive a document that appears to be from <ceo name>@yourcompany.com, they are far more likely to open it than if it comes from <randomname>@gmail.com.

45% OF ORGANISATIONS
HAVE AT LEAST ONE
COMPROMISED ACCOUNT
AND 6% HAVE AT LEAST ONE
VIP ACCOUNT COMPROMISED.

Impostor attacks continue to evolve. Figures 5-8 suggest a number of seasonal and industry trends, with threat actors continuing to focus on so-called many-to-many attacks. While early BEC attacks usually involved a single spoofed identity targeting a small number of users, attackers now often spoof many identities and send impostor attacks to many individuals within a given organisation. These attacks tend to be much smaller in scale than traditional malware attacks that can blanket an organisation, but the broader targeting appears to be netting greater returns for actors than earlier approaches.

The attacks are also tailored to the targeted industry, with education, for example, favouring subjects related to “Requests” and “Greetings” while engineering firms are more likely to see “Urgent” subject lines. Seasonal trends, on the other hand, show payment demands, for example, accelerating in the first quarter of 2018 and 2019. Yet company size continues to have no appreciable effect on the frequency or nature of impostor attacks. Social engineering in these attacks is effective and lucrative, regardless of company size.

This type of tailor-made attack represents the current state of the art in social engineering and exploitation of the human factor.

CONCLUSION

The human factor defines the actions and motivations of most threat actors today. The vast majority of attacks prey on people and human nature. Even automated exploits are frequently deployed in ways that still require initial execution by users rather than devices. With the widespread deployment of SaaS platforms and rising incidence and sophistication of impostor attacks, the stakes are higher than ever for organisations. People remain the primary target of attackers and the last line of defence for organisations, making a focus on people, as well as more traditional layers of security and training, critical to a holistic approach to defence.

RECOMMENDATIONS

Today's threats require a people-centric approach to keeping users safe. We recommend the following as a starting point:

- **Adopt a people-centred security posture.** Attackers do not view the world in terms of a network diagram. Deploy a solution that gives you visibility into who's being attacked, how they're being attacked and whether they clicked. Consider the individual risk each user represents, including how they're targeted, what data they have access to and whether they tend to fall prey to attacks.
- **Train users to spot and report malicious email.** Regular training and simulated attacks can stop many attacks and help identify people who are especially vulnerable. The best simulations mimic real-world attack techniques. Look for solutions that tie into current trends and the latest threat intelligence.
- **At the same time, assume that users will eventually click some threats.** Attackers will always find new ways to exploit human nature. Find a solution that spots and blocks inbound email threats targeting employees before they reach the inbox. Isolate suspicious and unverified URLs in email. And stop outside threats that use your domain to target customers.
- **Build a robust email fraud defence.** Email fraud can be hard to detect with conventional security tools. Invest in a solution can manage email based on custom quarantine and blocking policies.
- **Protect your brand reputation and customers in channels you don't own.** Fight attacks that target your customers over social media, email, and the web—especially fake accounts that piggyback on your brand. Isolate users' personal browsing and webmail from your environment. And look for a complete social media security solution that scans all social networks and reports fraudulent activity.
- **Partner with a threat intelligence vendor.** Focused, targeted attacks call for advanced threat intelligence. Leverage a solution that combines static and dynamic techniques to detect new attack tools, tactics and targets—and then learns from them.



For the latest threat research and guidance about
today's advanced threats and digital risks, visit
proofpoint.com/uk/threat-insight

ABOUT PROOFPOINT, INC.

Proofpoint, Inc. (PFPT) is a leading cybersecurity company that protects organisations' greatest assets and biggest risks: their people. With an integrated suite of cloud-based solutions, Proofpoint helps companies around the world stop targeted threats, safeguard their data and make their users more resilient against cyber attacks. Leading organisations of all sizes, including more than half of the Fortune 1000, rely on Proofpoint's people-centric security and compliance solutions to mitigate their most critical security and compliance risks across email, the cloud, social media and the web. More information is available at www.proofpoint.com.

©Proofpoint, Inc. Proofpoint is a trade mark of Proofpoint, Inc. in the United States and other countries. All other trade marks contained herein are property of their respective owners.