


REPORT

proofpoint.

THE HUMAN FACTOR

2017

proofpoint.com



The headline story of 2016 threat landscape was the explosive growth of ransomware and the massive email campaigns that delivered it to organizations of all sizes around the world. These attacks added up to billions of dollars in direct financial losses.

Cyber criminals relied less on automated attacks and exploits, shifting instead to social engineering. The change increased the impact and effectiveness of these campaigns. From email to software as a service, from social media to mobile apps, cyber criminals carried out social engineering at scale. They combined sophisticated, targeted lures and persuasive tricks with broad distribution. They employed new and improved techniques.

The tactics worked. Attackers tricked people into installing malware, handing over their credentials, disclosing sensitive information and transferring funds.




TABLE OF CONTENTS

Key Findings	4
By the Numbers	7
Human-driven exploits and thefts over binaries	7
Malware targeting trends	8
Malicious message volume by day of week	9
Thursday isn't just for throwbacks: Malware sending trends by category	10
Lures too good to resist	11
Phishing scale and effectiveness	11
Comparing effectiveness of small and large-scale phishing lures	12
Clicking behavior: Clicking from a work PC is so 2014	13
Physical industries, cyber risks	13
Working lunches are clicking lunches.....	14
Time to click	15
Automating the Human Exploit	15
Business email compromise (BEC): Exploiting the human factor.....	15
Spear-phishing at scale: Mass personalization automates social engineering	16
This time it's personal—anatomy of a personalized attack.....	17
Social engineering goes mobile.....	18
Social media phishing: The killer app.....	19
Fraudulent mobile apps: Exploiting the human factor on the go	20
Conclusion & Recommendations	22



Attacks exploit people rather than code.

Accelerating a shift that began in 2015, cyber criminals aggressively adopted attacks in 2016 that relied on clicks by humans rather than exploits of vulnerable software. By December, more than 99% of attachment-based email attacks were enabled by the user clicking something rather than an automated exploit. This trend extended to URL-based threats, where more than 90% of messages led users to credential phishing pages, which trick victims into entering their usernames and passwords, rather than to exploits.

Recommendation: Deploy solutions that can detect malicious macros and other code embedded within attachments designed to trick users into running them. Adopt and deploy solutions that can perform proactive and real-time sandboxing of URLs and attachments. By analyzing what happens if someone clicks on the link or attachments, sandboxing can reveal malicious behavior that's not obvious using conventional defenses. Your solution should also be able to recognize websites designed to steal credentials, even if the site is not known to be malicious.

KEY FINDINGS

Highly personalized, targeted email campaigns focus on exploiting people, not just their technology.

Spear-phishing email campaigns, which target specific people rather than indiscriminately seeking victims, were automated to operate at scale. Despite their large numbers, many included multiple personal details specific to the targeted recipient. Social engineering campaigns used documents with malicious macros and other techniques that tricked users into installing malware.

Recommendation: Deploy solutions that can detect and block sophisticated phishing messages before they reach the intended targets.

Mobile threats eschew exploits and use fraudulent mobile apps and next-generation SMS phishing to target customers of major banks and other consumer brands.

Attackers mimicked trusted brands, published apps with misleading names, and employed other ruses to convince users to download malware on their mobile devices. Users willingly downloaded and installed fraudulent apps that steal personal information and in some cases can take full control of mobile devices. SMS phishing, in which attackers use text messages to trick users into providing login credentials and other sensitive information, also increased. This trend reflects attackers' growing efforts to target users on devices they use the most, circumventing established network- and PC-based defenses.

Recommendation: Adopt and deploy mobile security solutions that work for company- and employee-owned devices. These solutions should be able to detect and stop next-generation SMS phishing attacks; detect, track, and block user clicks; and detect the presence of fraudulent, risky, and malicious apps on smartphones and tablets. At the same time, banks, telecom companies, retailers, and other organizations should adopt solutions that enable them to detect apps that misuse their brand to target their customers for theft, fraud, and other forms of abuse.

Social media fraudulent support account phishing increased 150% in 2016.

"Angler phishing" attacked customers of banks, social media, and other services using targeted responses to customer posts on brands' legitimate social media channels. Angler phishing is a term we use to describe attacks in which the attacker creates a lookalike social-media account posing as the customer-service



BEC attacks catch up to banking Trojans as attackers use people rather than binaries to steal funds.

As a share of the total volume for all email-based financial fraud, business email compromise (BEC) attacks exploded as attackers shifted away from techniques that rely on malicious programs to those that trick victims into carrying out the attack themselves. BEC attacks, where attackers trick victims into wiring money or sending sensitive information by posing as a colleague, accounted for 1% of email-based financial fraud messages in 2015, a far smaller share than banking Trojans. By the end of 2016, BEC attacks accounted for 42%. This increase was enabled by technical innovations such as large-scale subdomain spoofing and changes in how attackers target recipients and spoof senders.

Recommendation: Deploy a solution that can classify email dynamically as attacks change. Use this solution to build quarantine and blocking policies to stop attacks such as BEC, which are highly-targeted, arrive in low volumes at targeted organizations, and often have no payload at all—and are thus difficult to detect.

account of a trusted brand. When someone tweets to a company looking for help, the attacker swoops in. Victims are often directed to realistic-looking landing pages and tricked into handing over their account credentials.

Recommendation: Protect your brand reputation and customers. Fight attacks targeting your customers over social media, email, and mobile—especially fraudulent accounts that piggyback on your brand. Look for a robust social media security solution that scans all social networks and reports fraudulent activity.

Half of the clicks on malicious URLs occur on devices that are outside the purview of enterprise desktop management.

Some 42% of clicks on malicious URLs are made from mobile devices—more than doubling the long-running rate of 20%. And 8% of clicks occur on potentially vulnerable versions of Windows for which security patches are no longer available.

Recommendation: Adopt solutions that can protect employees from email-based attacks regardless of where they may end up reading the messages. Your solution should detect and block clicks on malicious URLs from smartphones, tablets, and email accesses through websites on a PC. For users who are still working from Windows PCs—whether company-owned or personal—ensure that all of them are on versions of Windows that are still supported with security patches. And be sure that all available operating system and application patches are installed.

Physical industries, cyber risks

Across all industries, the average click rate of 4.6% means that users click almost 1 in 20 malicious URLs. The highest click rates are concentrated in physical, old-economy industries such as mining and construction (“moving atoms”) rather than digital-era industries that handle personal, financial, and healthcare data (“moving bits”). All organizations and industries are targeted by modern cyber criminals, not just those most enmeshed in the digital economy.

Recommendation: Organizations in industries with large numbers of non-office employees must protect employees as aggressively as their peers do in financial services, healthcare, and technology. They must adopt and deploy enterprise-wide solutions that stop latest email-based attacks.



Thursday isn't just for throwbacks: malware categories vary distribution by day of the week.

Malware delivery times tend to be consistent every week, though with crucial differences between malware types and delivery vectors. Campaigns that use malicious attachments arrive at the beginning of the business day and drop off sharply after 4-5 hours; those that use malicious URLs arrive more evenly throughout the day. Threat actors time message delivery to maximize their impact. Information stealers arrive early in the week when they can collect the most information. Ransomware and point-of-sale (POS) Trojans arrive later in the week when security teams have less time to detect and mitigate infections before the weekend.

Recommendation: Organizations should increase monitoring for the presence of malware in their environment in the second half of the business week. And they should deploy automated incident response solutions that enable them to quickly resolve security incidents and mitigate threats in hours rather than days.

Almost 90% of clicks on malicious URLs occur within 24 hours after they're delivered.

These messages have their greatest impact the day they arrive: 87% of clicks occur within first 24 hours of delivery. Almost half of clicks occur within an hour after the message arrived. And a quarter of clicks occur just 10 minutes after arrival. The median time-to-click (the time between arrival and click) is shortest during business hours: from 8 a.m. to 3 p.m. EDT in the U.S. and Canada, the median time-to-click is less than 1 hour, a pattern that generally holds for the U.K. and Europe as well.

Recommendation: Quickly detecting malicious messages that are delivered and clicked is vital to reducing their potential impact. Organizations should deploy solutions that can proactively flag already-delivered messages and block clicked URLs found to be malicious after delivery. The longer a malicious URL sits in a recipient's email inbox, the more likely it is to be clicked. The first 24 hours in particular are critical to limiting the risk from delivered threats.

Working lunch? Clicking lunch.

Attackers understand when recipients are most likely to click on malicious messages and optimize their campaigns. Activity increases quickly with the start of the business day and peaks around 4-5 hours after that—right around lunchtime.


This pattern is largely consistent across other regions. Users in the U.S., Canada, and Australia follow this trend most closely, while French clicking peaks around 1 p.m. On the other hand, Swiss and German users don't wait for lunch to click; their clicks peak in the first hours of the working day. U.K. workers pace their clicking evenly over the course of the day, with a clear drop in activity after 2 p.m.

Recommendation: Deploy solutions that can protect users regardless of where they are reading and clicking on malicious messages—whether that's at their desks in the morning or on their smartphones over lunch.

BY THE NUMBERS

HUMAN-DRIVEN EXPLOITS AND THEFTS OVER BINARIES

By the second half of 2016, the shift to human-driven exploits was well-established. A full 99% of email-based financial fraud attacks relied on human clicks rather than automated exploits to install malware.



This trend also extended to URL-based malicious messages. On average, 90% of malicious URL messages per month led to credential phishing pages (sites designed to look like official login pages to trick users into providing account credentials), rather than to exploit kits (sites set up to detect and exploit vulnerabilities of machines connecting to it).

EXPLOIT KITS EMBRACE THE HUMAN FACTOR

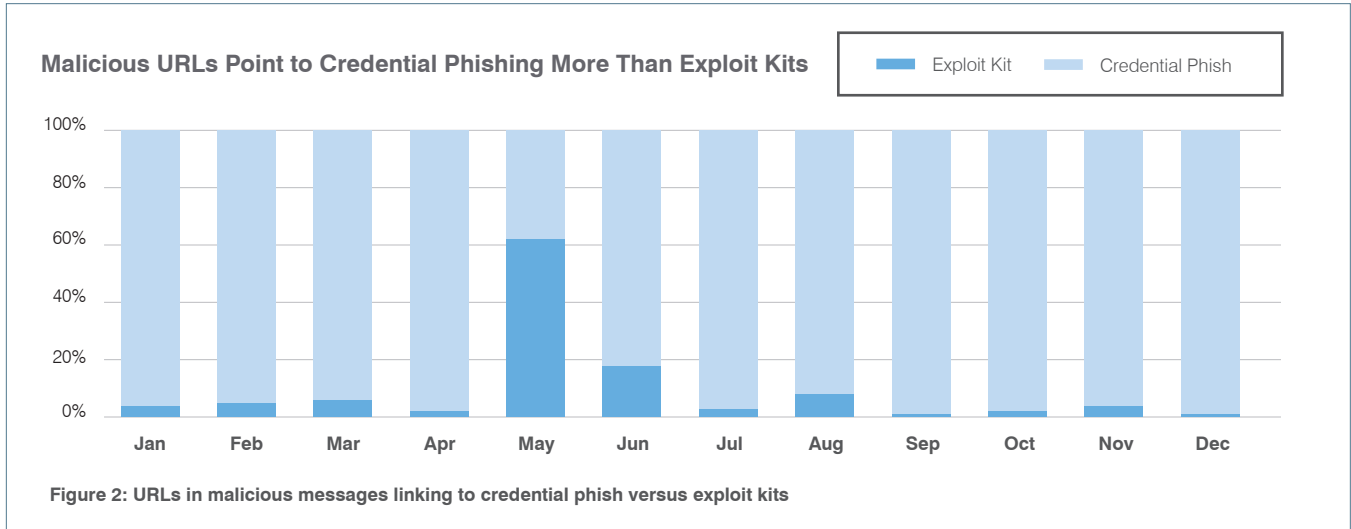


Few attack vectors would seem more bound to automated software exploits than exploit kits. But even here we see attackers shifting techniques to embrace the human over the exploit.

An established player in the exploit kit (EK) market Magnitude EK has operated for several years. It is fed by malvertising (online ads that hide malicious code to snare people visiting legitimate websites). Magnitude filters and redirects traffic so that only select targets are affected, most recently distributing Cerber almost exclusively in Korea and Taiwan. But in the first quarter of 2016, Magnitude began using a new social engineering chain affecting Internet Explorer users on Windows 10.

In this attack, malvertising on a legitimate web site directs targeted web surfers to a landing page. The page uses code that prevents users from closing or bypassing dialog boxes. A series of on-screen prompts that leverage expected Windows dialogs and behavior lead users to download a shortcut containing Windows PowerShell commands. The command, in turn, downloads and executes Cerber ransomware.

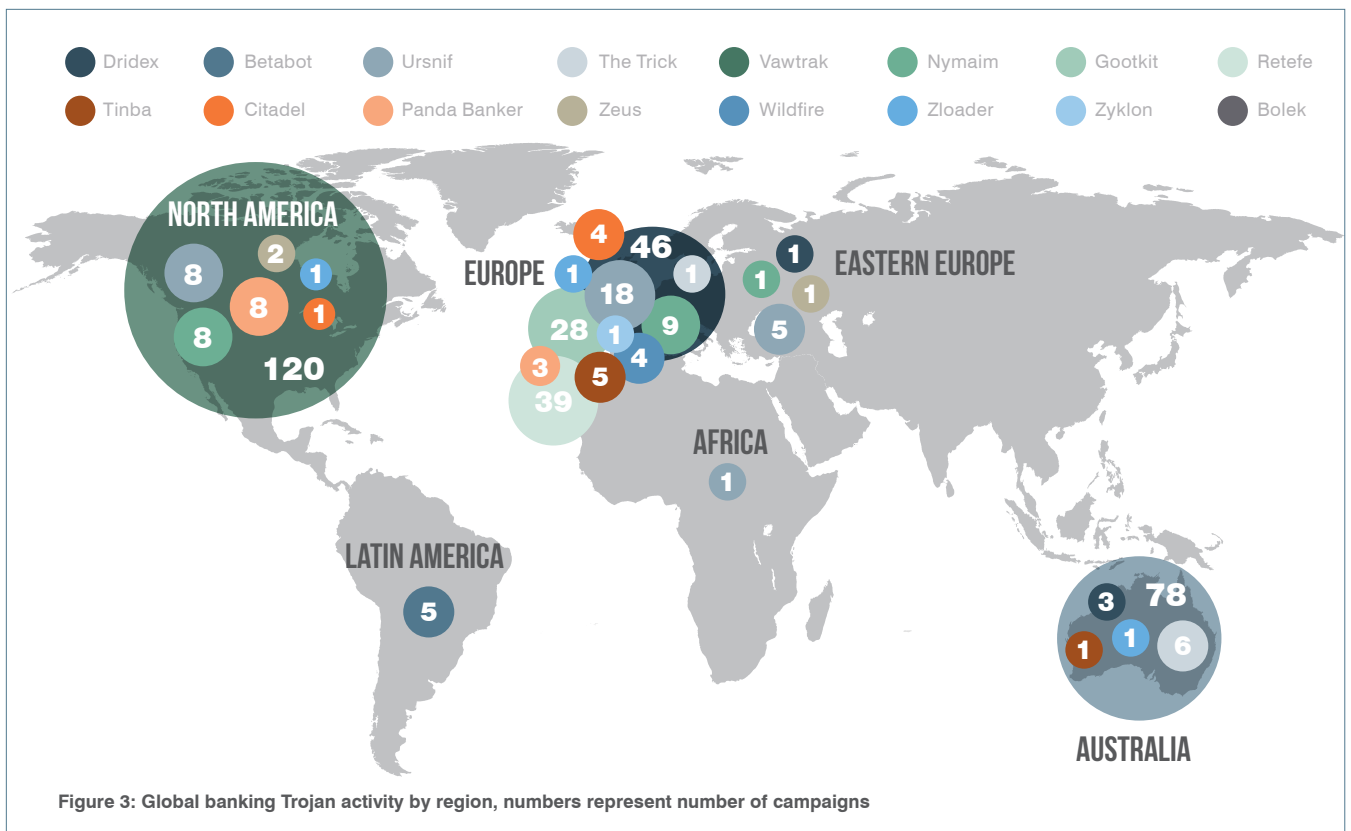
This social-engineering scheme lacks the refinement of some email-based malware campaigns, but it underscores the breadth of attackers' shift to exploiting "the human factor." Social engineering-based approaches free the exploit kit from the constraints of automated software exploits. Instead, these attacks entice users to click buttons, bypass sandboxes, run PowerShell code, and perform other actions that infect their own systems.



Between them, these two changes epitomize the focus by cyber criminals on attack techniques that leverage human interaction rather than automated exploits to infect systems, steal credentials, and transfer funds.

MALWARE TARGETING TRENDS

Attackers used banking Trojans to target victims in specific geographies, further echoing the trend of human-driven exploits. These attacks use lures and attachments in local languages; region-specific web code that relays malicious instructions (web injects); and campaigns timed to align with the business day – and clicking behavior – of their intended recipients.

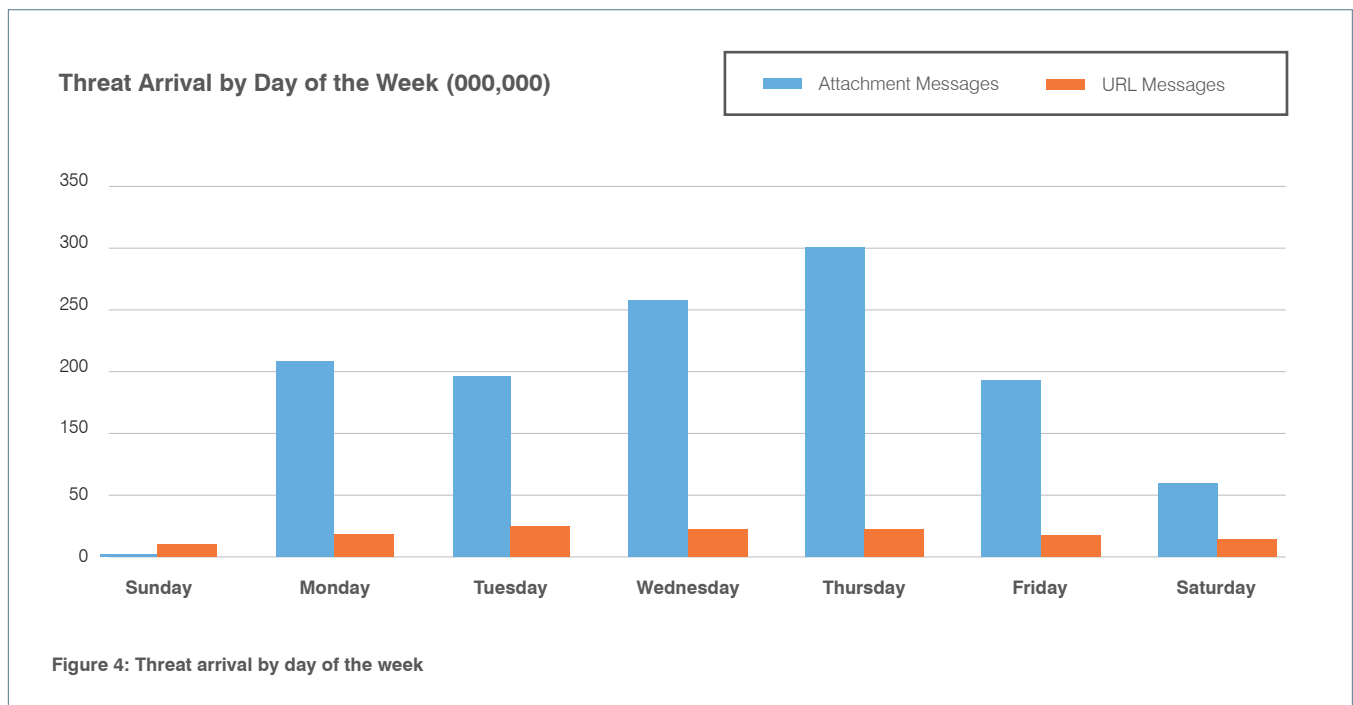


Recommendation: Deploy best-of-breed defenses that can detect and block regionally targeted campaigns. Get threat intelligence that keeps your team informed about the latest trends in attackers and how they target payloads.

MALICIOUS MESSAGE VOLUME BY DAY OF WEEK

Email-based threats can target users any day of the week, and attackers optimize the days and times of their campaigns for the biggest impact. Attackers do their best to make sure messages reach users when they are most likely to click: at the start of the business day in time for them to see and click on malicious messages during working hours.

Email-based threats arrive every day of the week, but some days bring more attacks than others. Figure 4 shows the typical weekly pattern:



Malicious attachment message volume spikes more than 38% on Thursdays over the average weekday volume. These weekday targeting trends appear to be global: Thursday is the top message volume day for attachments in all the countries we examined.

The numbers behind these trends reveal some important patterns:


- Malicious URL message volume is more evenly distributed across the weekdays. Tuesday and Thursday remain the top days for sending malicious URL messages – the main vector for credential phishing attacks – although the variance is less pronounced than in previous years.
- Weekends are still low-volume days for email-borne threats. Still, URL message volume does not drop off as significantly as attachments—especially compared to previous years when weekend message traffic was negligible.

We also see more regional variation for URL campaigns by day of the week. Of the regions we examined, sending days for Europe diverged the most from the U.S. and Canada. Thursday is the clear peak day, accounting for 20.2% of weekday message volume. Tuesday is next at 17.6%. Monday, Wednesday, and Friday are roughly equal at about 15% each.

Recommendation: Adopt defensive solutions that can protect your users from the full-range of email-based attacks seven days a week. The solutions should have the capacity to handle the highest message volume days without impeding performance or sacrificing effectiveness.


THURSDAY ISN'T JUST FOR THROWBACKS: MALWARE SENDING TRENDS BY CATEGORY

Malware campaigns are not evenly distributed across the week. Instead, they exhibit clear patterns, with some malware categories favoring some days of the week over others.



Ransomware, which locks away victims' data until they pay a fee to unlock it, is one example. Ransomware message volumes were much higher on Thursday than the other days of the week, driven primarily by high-volume campaigns that sent the Locky strain. With few exceptions, ransomware was the only category of malware sent on weekends.

The ransomware campaigns of 2016 were large, with volumes that obscure trends in other malware categories. Figure 6 shows message volume with ransomware excluded.



While Wednesday is the peak day for banking Trojans, credential stealer campaigners favor Thursdays, and downloaders are spread relatively evenly across Monday through Thursday.

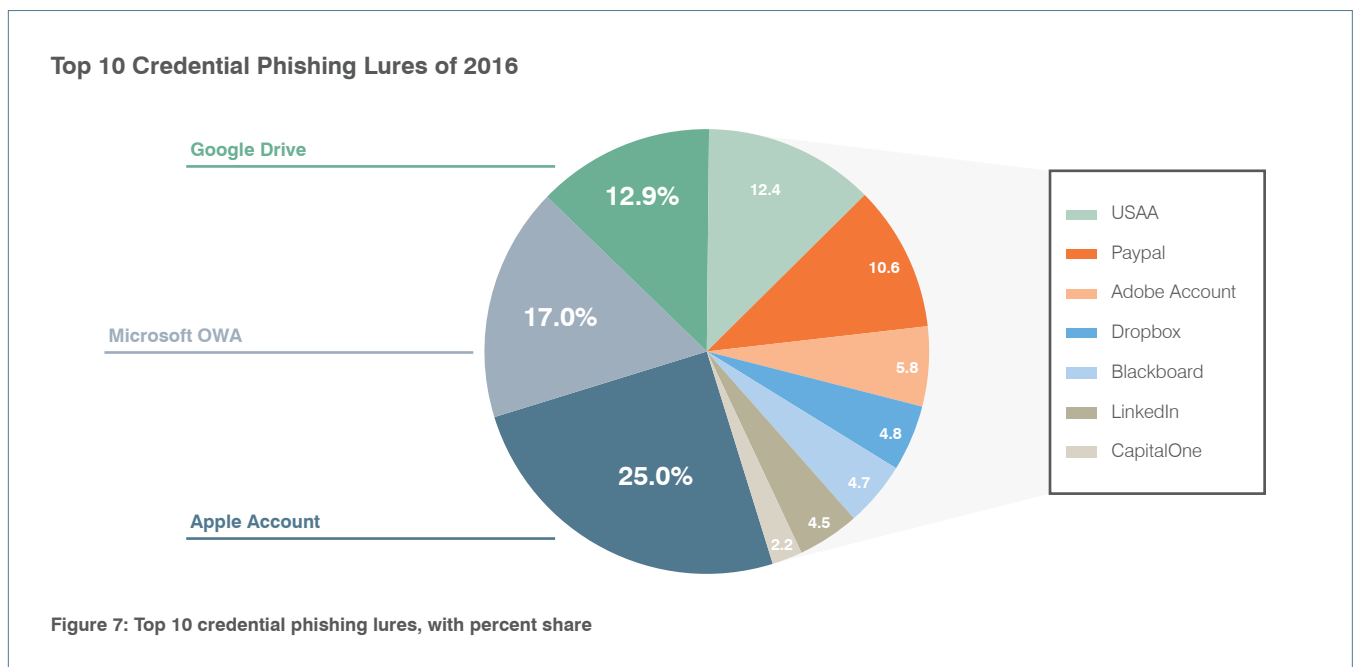
The numbers for lower-volume malware show even clearer preferences in sending days:

- Keyloggers and backdoors favor Mondays. The number of Monday campaigns for backdoors was 68% greater than the Tuesday-through-Thursday average. The Monday bias of keyloggers was even more pronounced; more than twice as many keylogger campaigns send on Mondays than the Tuesday-through-Thursday average.
- Point-of-sale (POS) campaigns were sent almost exclusively on Thursdays or Fridays, with 80% of 2016 campaigns occurring on one of those two days.

LURES TOO GOOD TO RESIST

Phishing scale and effectiveness

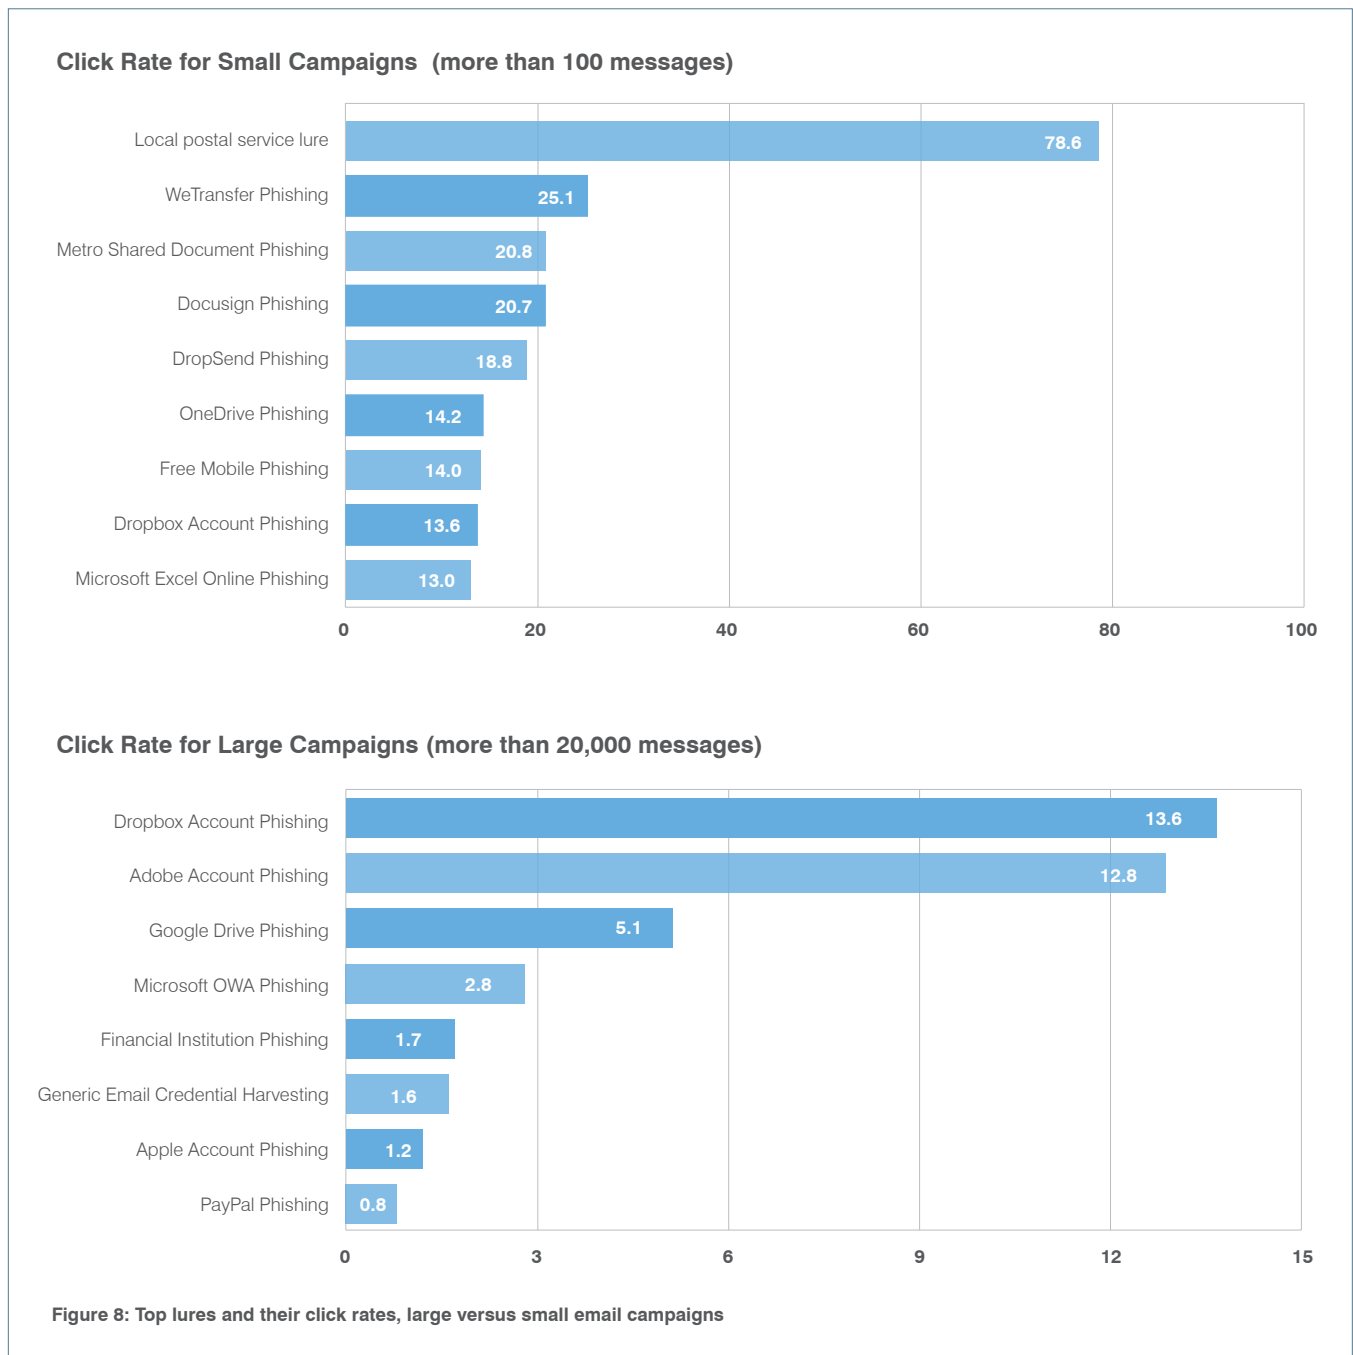
The most popular credential phishing lures were consistent with those in 2015. In fact, the five most common lures were almost unchanged.



As was the case the year before, 2016 delivery volume did not correlate to click rates. Phishing messages designed to steal Apple ID were the most sent, for example, but Google Drive phishing links were the most clicked.

Accounts used to share files and images—such as Google Drive, Adobe Creative Cloud, and Dropbox—are the most effective lures. These messages made up less than 24% of the message volume among the top ten lures but were the most effective as measured by click rates. We also saw a pronounced difference in the lures and effectiveness between the large and small credential phishing campaigns, as shown in Figure 8.

Comparing effectiveness of lures by campaign size

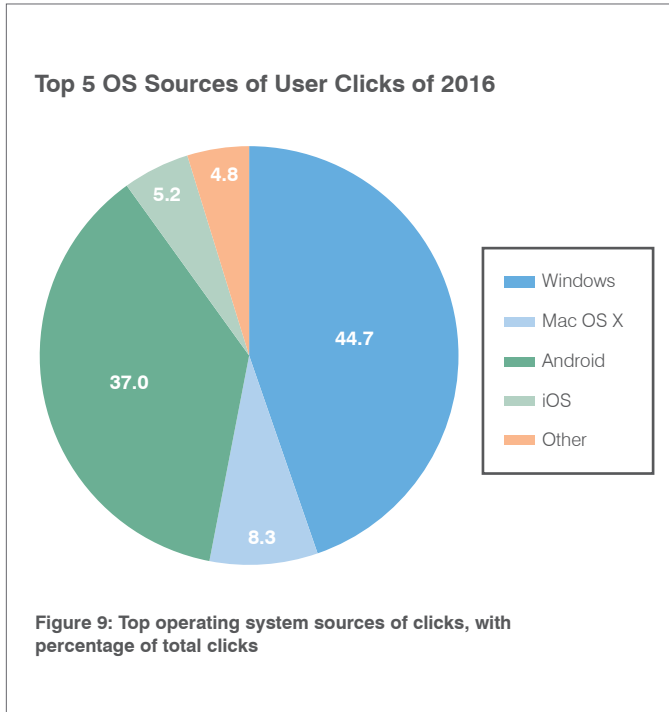


While social-media lures don't make a huge statistical dent in our threat data, they can still be effective in smaller, more targeted campaigns. Document sharing lures, meanwhile, are consistently effective (and thus popular with attackers) in large and small campaigns. More important, smaller campaigns drive a higher click rate than large campaigns. That makes quickly detecting and mitigating them crucial.

Recommendation: Teaching employees to beware the latest and most effective phishing lures is important. But attackers can change lures, payloads, and any other aspect of their campaigns overnight. Deploy solutions that can detect a variety of credential phishing attacks through a combination of proactive and real-time URL sandboxing in emails.

CLICKING BEHAVIOR: CLICKING FROM A WORK PC IS SO 2014

In 2014, 91% of user clicks occurred from Microsoft Windows PCs. In the last two years, that percentage has fallen by half. Over the same period, the percentage of clicks from mobile devices more than doubled, to 42% of total clicks on malicious URLs.



Users have shifted their work—and their clicking—to mobile devices. Attackers are taking advantage of this shift to target users with social media, banking, and other mobile apps that can trick users into giving up sensitive information, no automated software exploit needed.

But clicks from Windows PCs remain a problem for security managers, as is clear from the following:

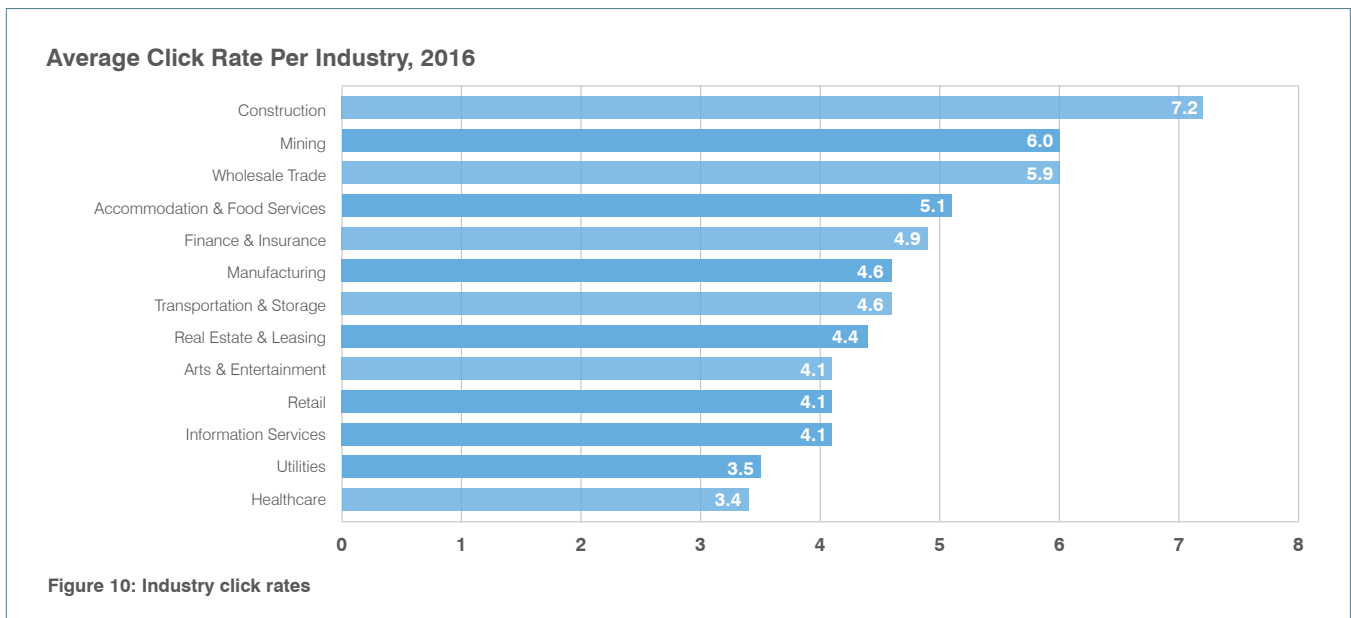
- 66% of Windows PC clicks (29% of total clicks) occur on a version of Windows that is no longer in “mainstream support” (Windows 7)
- 19% of Windows PC clicks (8.5% of total) are from versions that are no longer issued security patches at all (Windows XP, Windows Vista, Windows 2000), almost doubling in proportion compared to 2014.

While attackers are relying less and less on automated exploits, one of the most commonly used exploits in email attachment attacks takes advantage of a four-year-old Microsoft Office vulnerability (CVE-2012-0158). That’s why deploying operating system and application patches as quickly as possible is still important.

CLICK TRENDS BY INDUSTRY

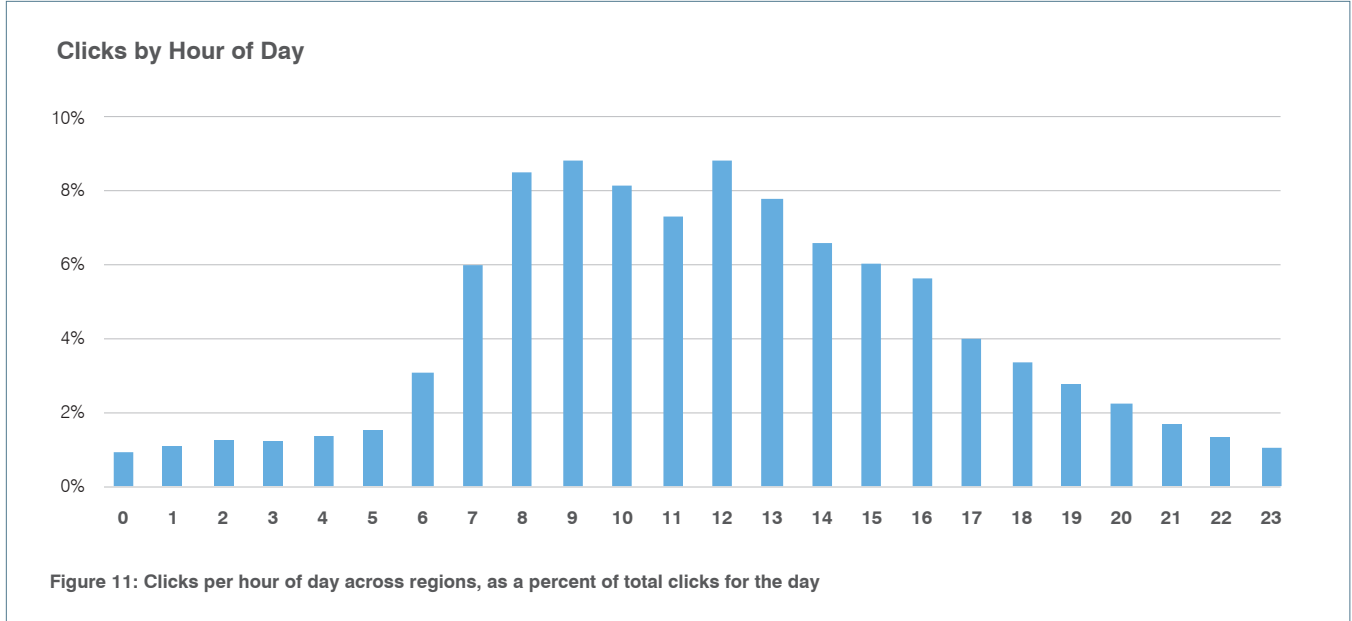
The 2016 data proved yet again that every organization clicks: 4.6% of malicious URLs are clicked across all industries and organizations. As in past years, some industries click more than others. Industries that “move atoms”—for example, construction and mining—click more often on malicious URLs than digital-era industries that “move bits.”

4.6% Average click rate on malicious URLs across all industries




WORKING LUNCH? CLICKING LUNCH

The times of day at which users click on malicious URLs are consistent across regions. Activity kicks off rapidly with the start of the business day, peaking 4-5 hours later—right around lunchtime.



The most striking takeaway from Figure 11 is that users click on malicious URLs at every hour of the day. Whether at work or at home, day or night, users are clicking on URLs that can lead to phishing pages and malware downloads.

Within the daily cycle, we see some regional variations in the times during which users are most likely to click on malicious URLs, as Figure 12 shows.




Click activity varies by country:

- Users in the U.S., Canada, and Australia peak midday. French clicking peaks around 1 p.m.
- Swiss and German users don't wait for lunch to click; clicks peak in the first hours of the working day.
- UK workers click evenly over the course of the day, with a clear drop in activity after 2 p.m.

TIME TO CLICK

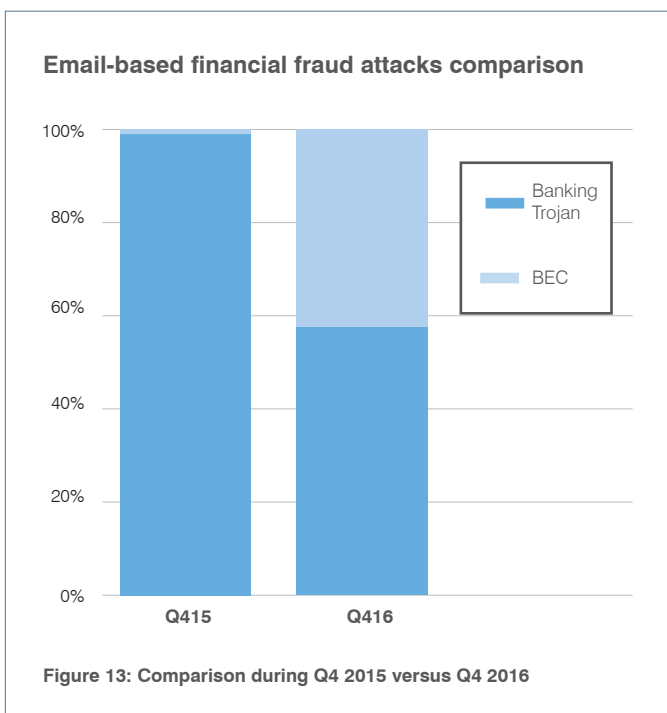
The peak clicking times coincide with business hours. During this time, malicious URLs are likely to have their shortest wait times before being clicked. The median time-to-click for malicious URLs is less than one hour during business hours.

Most clicks occur within one day after malicious URLs arrive in the user's inbox. Here's a breakdown of how quickly URLs are clicked after delivery:



AUTOMATING THE HUMAN EXPLOIT

BUSINESS EMAIL COMPROMISE (BEC): EXPLOITING THE HUMAN FACTOR



The rise of business email compromise (BEC) attacks highlights the growth of attack techniques that shift the burden of action from an automated exploit or tool to a human. Three quarters of our worldwide customer base experienced at least one BEC attack attempt in the last three months of 2016. This growth was reflected in the rise of BEC relative to banking Trojans in financial fraud attacks between 2015 and 2016.

BEC is a new type of threat, but attackers are already evolving their techniques in the face of increased user awareness and automated defenses. BEC attackers often would send spoofed messages to the CFO of a targeted company, purportedly from the CEO. That began to change in the latter part of 2016, as Figure 14 shows.

45%

Quarter-over-quarter increase in BEC attacks in Q4

Percent of BEC Emails Spoofed as “CEO to CFO”
July-December 2016




Figure 14: Percent of total BEC emails spoofed from the CEO and sent specifically to the CFO, June-December 2016

While CEO impersonation, or spoofing, continues in BEC attacks, cyber criminals are increasingly targeting victims deeper within organizations. Attacks are shifting beyond the CEO-CFO relationship, targeting the CEO’s relationship with other employee groups. They might target accounts payable for wire transfer fraud, engineering to steal intellectual property, and human resources to get confidential tax and identity information.

SPEAR-PHISHING AT SCALE: MASS PERSONALIZATION AUTOMATES SOCIAL ENGINEERING

Before 2016, email threat actors largely had to choose between two approaches:

- Large-scale “spray-and-pray” campaigns that sent hundreds of thousands or millions of malicious email messages to unfiltered recipients
- Small, highly targeted campaigns with carefully crafted lures

In 2016, however, a prolific actor we refer to as TA530 began distributing personalized emails in large-scale targeted campaigns that used sophisticated social engineering techniques. These campaigns often involved thousands or tens of thousands of messages targeted by industry vertical with malicious payloads geared towards that industry. For example, point-of-sale malware might appear in retail-targeted campaigns; banking Trojans and information stealers were often used against manufacturing or technology targets.

TA530 took advantage of data harvested from public sources such as LinkedIn and data from compromised online customer relationship management (CRM) systems. The actor used the information to develop highly personalized email lures and document attachments that featured recipient names, roles, addresses, company names, and more. These elements add a patina of legitimacy to the emails; social engineering adds urgency.

In one retail-targeted campaign, a bogus customer complaint referenced specific store addresses. The email included document attachments that supposedly provided more details. If the store didn’t address the issue promptly, the email warned, the “customer” would escalate the complaint. The attached documents contained malicious macros: clicking the Enable Content button installed AbaddonPOS, a point-of-sale malware variant.

In another campaign, recipients were told that attached documents were subpoenas for a court appearance. Again, social engineering and personal details in the email goad the user into opening the attachment. The opened document looks like a subpoena but also has social engineering features—an image overlay explaining why readers must enable macros to read the document. Enabling the macros downloads the Ursnif banking Trojan onto the victim’s PC.

TA530 distributed a wide range of malware for other cyber criminals. The high degree of personalization and clever social engineering present in their campaigns exploited the human factor at scale. That scale exposed large numbers of users to banking Trojans, information stealers, and more by tricking even savvy users into opening attachments and running malicious code.

THIS TIME IT'S PERSONAL—ANATOMY OF A PERSONALIZED ATTACK

In addition to the common tricks attackers use to catch the attention of a potential victim and create a sense of urgency around the message, the personalized emails used by TA530 to distribute malware incorporated details specific to the recipient's company. Previously common in small-scale, hand-crafted spear phishing attacks, the TA530 campaigns of 2016 automated these techniques to carry out large-scale social engineering through email-based attacks.

The image shows a screenshot of an email client window displaying a phishing message. The email is titled "Appearance Notice AP1691788" and is from "Meghan Cox <loisbaker@cox.net>". The body of the email contains a hearing notice for a case on April 2nd, 2016, at 11:00 AM, before Justice Sandra Peterson. The email includes an attachment named ".doc (114 KB)". The sender's name is "Meghan Cox" and the company name is "Hyland, Mark J. Attorney". The email also includes a complaint reference number "A01047597".

Annotations in orange boxes point to specific details in the email:

- Subject line includes the name of the recipient business**: Points to the subject line "Appearance Notice AP1691788".
- Company name appears in name of fake legal complaint**: Points to the subject line "Appearance Notice AP1691788".
- Attachment name is name of targeted company**: Points to the attachment ".doc (114 KB)".
- Name and valid street address of targeted company**: Points to the recipient's name and address in the "Attention: Owner / Executive at" line.
- This is a hearing about**: Points to the text "This is a hearing about" in the body of the email.

SOCIAL ENGINEERING GOES MOBILE

Though not new, SMS phishing that targets consumers and enterprises is on the rise, and actors are introducing new techniques to increase its effectiveness. Because there are no commercially available SMS inbound filtering products as there are with email, attackers have discovered that sending SMS messages can be highly effective for tricking users into handing over their banking credentials.

This defensive gap is compounded by the fact that the small screens of mobile devices make it difficult to determine whether websites are fake. In the past, SMS phishing usually involved a text message with a single link to a fake account login page, often for telecoms and other accounts. By late 2016, we saw attackers adding new techniques and twists to better leverage the potential effectiveness of SMS phishing schemes.

Consider the example of a set of late-2016 SMS-based phishing messages purporting to be from a major U.S. bank. The messages were received from email addresses and a phone number and bore legitimate-looking links.

Instead of taking users directly to a phishing form or site, the links shown in Figures 15 and 16 first take users to an image (Figure 17). This technique defeats many phishing filters because the brand name is an image and, therefore, cannot be parsed by automated tools. After six seconds, victims are automatically redirected to the real phishing site.

42% User clicks on malicious URLs from mobile devices

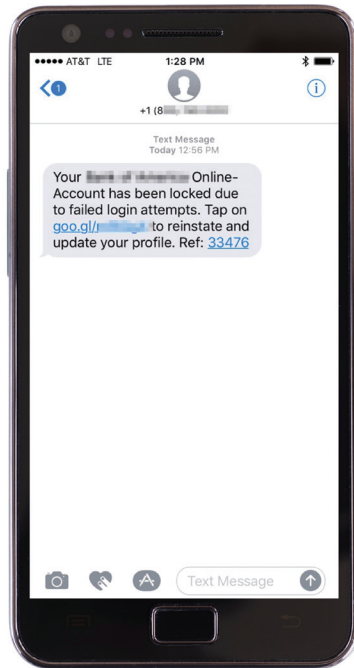


Figure 15: SMS phishing message originating from a phone number

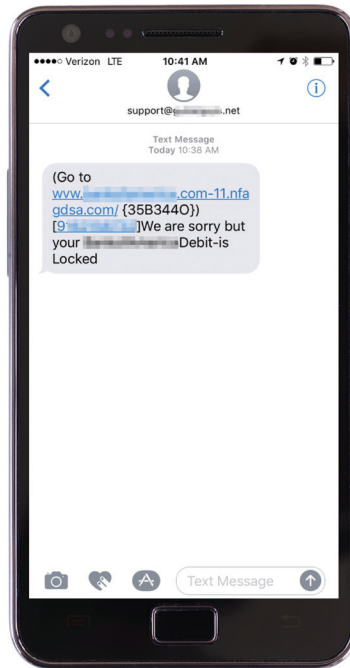


Figure 16: SMS phishing message originating from an email address


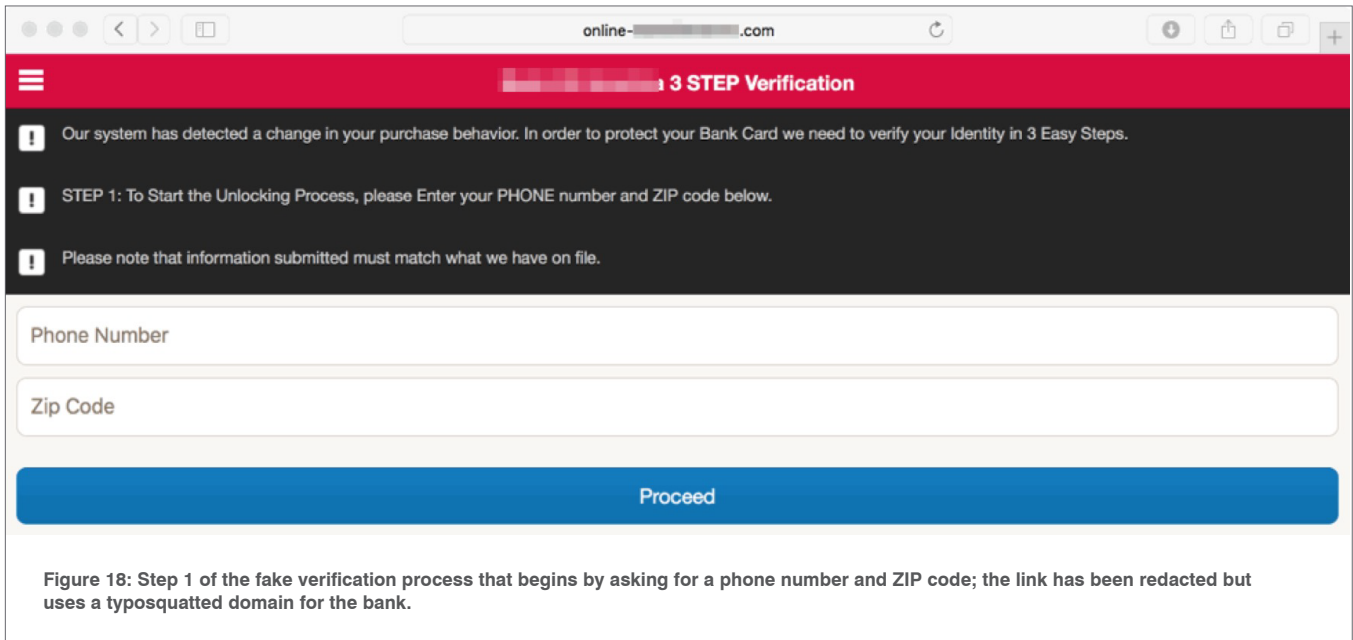


Figure 17: Image initially linked from phishing messages with stolen branding redacted

The phishers then present a clever three-step verification that begins with the victim's phone number and ZIP code instead of more "traditional" phishing sites that begin by asking for passwords and usernames.



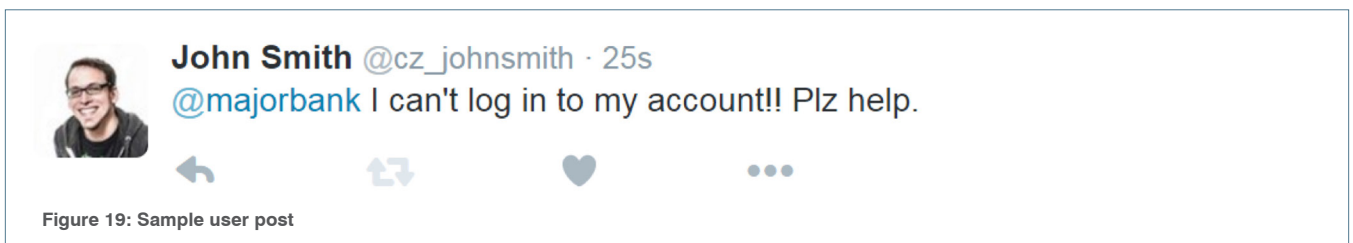
Victims are then prompted to enter an email address in the next step of the fake verification process. If recipients enter a Gmail or Yahoo address, they are presented with realistic but bogus login page for those services. If victims enter their password, the attackers can gain control of the Gmail or Yahoo account and can reset passwords for any other services attached to the email address.

By this stage of the phishing attempt, attackers have captured their victim's phone number, ZIP code, email address, and email password. The final stage of the attack brings victims back to the bank phishing site and prompts victims to enter credit card information and their Social Security number. Even if recipients become suspicious at this point, attackers already have a phone number and access to an associated email account. For many providers, this is enough data to port the phone number away from the original provider and take control of a victim's online identity. In many cases, recipients also enter credit card data and Social Security numbers, allowing the attackers to rack up credit card charges and steal victims' identities.

SOCIAL MEDIA PHISHING: THE KILLER APP

After emerging in late 2015, fraudulent social media customer service accounts became a major feature of the threat landscape in 2016. These fraudulent accounts, which impersonate popular brands and respond to customer requests, can appear legitimate. These so-called "angler phishing" attacks grew 150% in 2016. And over the course of the year, more brands and industries were phished.

Consider the example of Twitter user "John Smith," who tweets a question to his bank, illustrated here as "Major Bank" with the Twitter handle "@majorbank":



Because the mention of “@majorbank” is a public action on Twitter, angler phishers are instantly alerted to the fact that John Smith is a customer of Major Bank and that he is asking for help with his account. Armed with this information, attackers insert themselves into the conversation and respond to John Smith. Here’s an example of a typical reply:

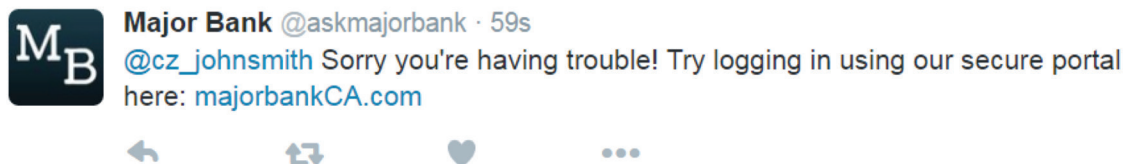


Figure 20: Sample reply as angler phishing attacker attempts to send original poster to phishing landing page

The fake account uses the logo of “Major Bank” and a handle (“@askmajorbank”) that sounds valid. Legitimate support accounts often have a different handle from the main Twitter account for the brand, and many brands have multiple accounts to serve different regions, product families and so on. So even if Smith noticed that the handle responding to him was different than the one he mentioned, it may not raise any suspicion. Social media threats are less common than email phishing, so most users are not aware that they are being attacked. If John Smith were to click on the “@askmajorbank” handle to see account profile page, most angler phish accounts use stolen branding to convincingly mimic a brand’s legitimate profile.

Happy to receive a prompt response and unwary of the responder’s authenticity, Smith clicks the link that takes him to “majorbankCA.com.” By all appearances, it’s Major Bank regular online banking login page. Smith enters his personal information, including his bank account number and login, to the fake login page. Smith has just been phished. Unfortunately, neither Smith nor Major Bank is aware that the account is now compromised.

When our researchers first discovered this form of phishing in late 2015, it mostly targeted the customers of big banks. These attempts appeared about 2-3 times per month, targeting a handful of accounts. By late 2016, the frequency had increased to 2-3 attempts per day for some major banks. And the attacks have expanded to target the customers of other industries, including online banks, media outlets, and entertainment companies. Many such attacks now include components for phishing email addresses and passwords, links to malware, and other criminal activity.



GOING AFTER GAMERS

Like many other targeted scams on social media, angler phishing attacks tend to cluster around the dates of upcoming compelling events. For instance, we saw angler phishing accounts enticing users to “win” in-game prizes by clicking on phishing links associated with a major gaming company. These types of accounts pose as legitimate accounts representing the games or the gaming company, and their activity spikes during launches of new gaming products.

FRAUDULENT MOBILE APPS: EXPLOITING THE HUMAN FACTOR ON THE GO

In 2015 and early 2016, cyber criminals often cloned popular games, adding malicious code that the user doesn’t see. By the end of 2016, mobile attacks targeted customers of specific banks, employees in particular industries, event attendees, and more. Attackers use stolen branding, misleading app names, and other ruses to convince users to download hidden malware on their mobile devices.

For example, we recently analyzed a [sample Android app](#) distributed in China that purports to be a point-of-sale (POS) control app for a major POS system manufacturer. The icon, shown below and written in Mandarin, references the manufacturer and shows one of its actual POS systems.

During installation, the app requested extensive permissions that didn't match what the app was supposed to do. As the permissions suggest, the app is actually a robust information stealer that runs persistently in the background.

Another example of a masquerading app appeared on the iOS app store late in 2016. It claimed to be an online banking guide, providing a variety of tips and tricks for the official app of a major U.S. bank. The app's listing in the iOS app store appears below.

Once installed, the app offers a feed of news articles and a link to "More Apps" (Figure 24). While the guide is not malware, the articles in the feed often lead to phishing sites while the linked apps are side-loaded adware.

This is not an isolated issue. More than 1% of worldwide app developers—that's about 16,000 publishers—are distributing malicious apps through mainstream and third-party app stores. Most masquerade as legitimate apps but are anything but.



Figure 21: Icon for fake POS management app

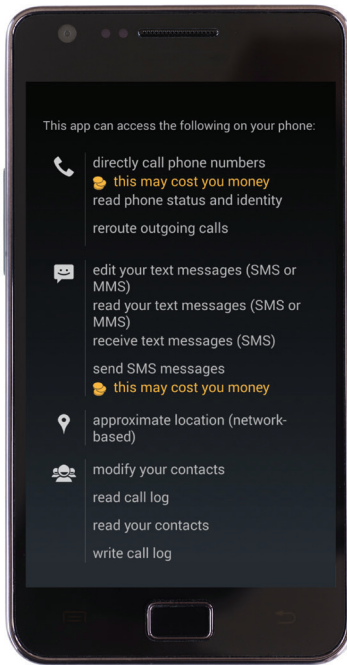


Figure 22: Part of the permissions requested by the bogus POS control app

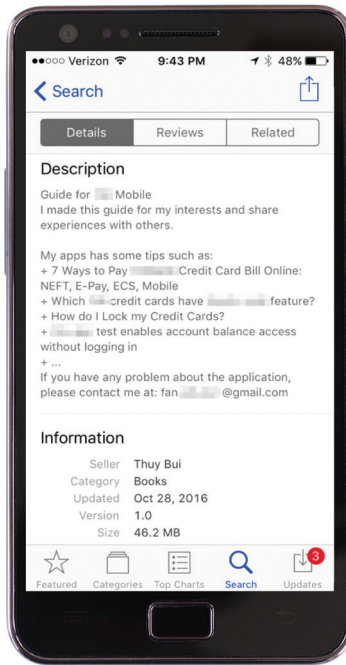


Figure 23: iOS listing for mobile banking guide

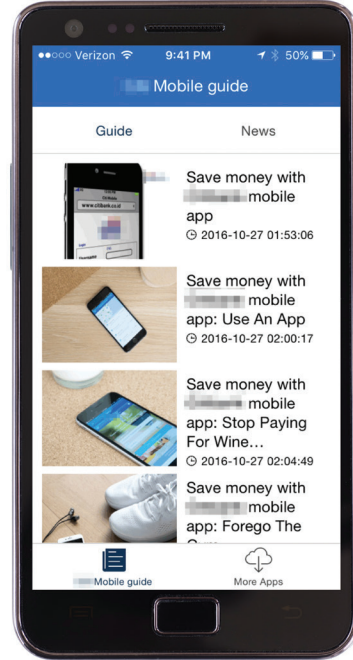


Figure 24: The guide app acts as a broker for side-loaded adware and links to a variety of phishing sites, some of which have been abused in other campaigns to deliver ransomware

CONCLUSION

Dramatic shifts in the threat landscape that started 2015 continued throughout 2016 and into 2017. Traffic for traditional exploit kits dropped by more than 94% in 2016, driven by factors ranging from law enforcement action to the increasing rarity of viable exploits. At the same time, ransomware exploded on the scene and targeted attacks grew to include new vectors used in tandem with email. Social media is now an integral part of an attacker's arsenal, with a 150% increase in angler phishing attacks as more brands and industries are targeted in these schemes.

Phishing campaigns moved to new channels in 2016, including mobile devices. Targets of these attacks will often receive SMS and email instructions asking for account credentials. Employees clicking on SMS messages with malicious links clicked 42% during 2016 compared to the long-running rate of 20%.

Human targeted attacks continued to lead the pack in 2016. Attackers' used automation and personalization to increase the volume and click-through rates of their campaigns. Taking a page from the B2B e-marketer's playbook, cyber criminals are adopting marketing best practices and sending their campaigns on Tuesdays and Thursdays when click-through rates are higher. Meanwhile, BEC and credential phishing attacks targeted the human factor directly--no technical exploits needed. Instead, they used social engineering to persuade victims into sending money, sensitive information and account credentials.

Timing is everything—attackers know that hitting your employees with a well-crafted email at the just the right time produces the best results. Of course, this varies by region. So if you are responsible for worldwide SecOps, you need visibility into not only attack patterns but also when and which employees tend to click.

RECOMMENDATIONS

Focus your security efforts on the leading vector for threats entering your organization: email. Deploy protection that works within the flow of email to stop attacks before they have a chance to reach your employees.


Detect threats in attachments and URLs with threat analysis services that use multiple approaches to examine behavior, code, and protocol. The earlier in the attack chain you detect malicious content, the easier it is to block, contain, and resolve.

Employ cloud-based sandbox analysis services that can scale to protect everyone in your organization. It should be able identify attack campaigns and uncover new attack tools, tactics, and targets so the next attack is easier to catch.

Protect employees in the field by providing same level of security controls to their mobile devices you provide for company-owned PCs in the office. Field workers are an increasing source of clicks on malicious links. And SMS messaging is emerging as a new attack vector. Your solution should detect and block clicks on malicious URLs on smartphones and tablets, on and off the network, regardless of location.

Fight fraud by stopping angler phishing or malicious apps trying to cash in on your brand. Get a solution that can scan and discover fraudulent accounts and applications that impersonate your company on social media and app stores.

In addition, accelerate your response to incidents. Consider a solution that enables you to retract malicious emails that have been delivered to users' inboxes. The solution should move malicious email out of users' hands and have the business logic to find and remove any copies of those messages that were forwarded.



ABOUT PROOFPOINT

Proofpoint, Inc. (NASDAQ:PFPT), a next-generation cybersecurity company, enables organizations to protect the way their people work today from advanced threats and compliance risks. Proofpoint helps cybersecurity professionals protect their users from the advanced attacks that target them (via email, mobile apps, and social media), protect the critical information people create, and equip their teams with the right intelligence and tools to respond quickly when things go wrong. Leading organizations of all sizes, including over 50 percent of the Fortune 100, rely on Proofpoint solutions, which are built for today's mobile and social-enabled IT environments and leverage both the power of the cloud and a big-data-driven analytics platform to combat modern advanced threats.

proofpoint.

www.proofpoint.com

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