2022 Social Engineering Report

KEY FINDINGS

- Threat actors may build trust with intended victims by holding extended conversations
- Threat actors expand abuse of effective tactics such as using trusted companies' services
- Threat actors leverage orthogonal technologies, such as the telephone, in their attack chain
- Threat actors know of and make use of existing conversation threads between colleagues
- Threat actors regularly leverage topical, timely, and socially relevant themes

WHAT WE KNOW

Overview

Social engineering is the preeminent component of the overwhelming majority of cyberattacks today. Whether the goal of a threat actor is to directly perpetrate fraud, harvest credentials, or install malware, at some point a human being must be coerced into taking an action on the actors' behalf. This fact is the basis for Proofpoint's People-Centric Security Model, an idea which has revolutionized the way the world's leading businesses consider the threat landscape and defend their organizations. It is the central driver of security awareness programs, which train end users to better recognize attempts to exploit them into facilitating malicious activity.

Despite defenders' best efforts, cybercriminals continue to defraud, extort, and ransom companies for billions of dollars annually. We are locked in an adversarial struggle with these threat actors, the nature of which evolves over time. As new defensive capabilities are implemented, crafty and technically talented actors look for new ways to defeat them. Security-focused decision makers have prioritized bolstering defenses around physical and cloud-based infrastructure which has led to human beings becoming the most reliable entry point for compromise. As a result, a wide array of content and techniques continue to be developed to exploit human behaviors and interests.

The most effective methods prey on natural human tendencies and undermine instincts which raise an alarm that "something isn't right." Often this means presenting the intended victim with content they may already be familiar with or regularly interact with in their day-to-day jobs: invoices, receipts, documents, and spreadsheets. The content appears routine and therefore raises no alarm. A threat actor might impersonate a trusted partner, or an authority figure such as a company's executive.

Social interest is also frequently leveraged: at the beginning of the COVID-19 pandemic there was a collective desire for information around updated health guidelines, company policies, regional mandates, and vaccine development. Because of the universal relevance, threat actors of every sophistication level pivoted to make use of COVID-19 related content.

Ultimately, the techniques which are successful will continue to be utilized and refined. What isn't effective will be discarded. The result is that throughout the year we observe threat actors constantly trialing new methods and content while improving those which already reliably earn clicks.

These developments coincide with the ability of the intended victims to recognize the threat actors' attempts as malicious. As users are better trained and become more aware, actors will be forced to pivot. In 2021, Proofpoint Threat Research noted social engineering content often aligned to key mistaken assumptions end-users continue to hold:

- The assumption that threat actors will not spend time building rapport prior to executing attacks, such as by holding regular conversations
- The assumption that legitimate services such as those provided by authoritative technology companies like Google and Microsoft are safe to use
- The assumption that threats only involve their computer and not orthogonal technologies such as the telephone
- The assumption that threat actors are unaware of email conversations held with colleagues and that those existing conversation threads are safe
- The assumption that threat actors won't make use of timely, topical, socially relevant content to pique interest or exploit emotions

This report provides evidence of how, throughout 2021, threat actors repeatedly subverted these assumptions to exploit the human element in their attacks.

Key Assumption: Threat actors don't have conversations with you

An important component of enticing people to interact with malicious content is to get them to trust the source. Effective social engineering is about generating feelings within a user that mentally drive them into engaging with content. Something is urgent, someone is trustworthy, someone can help. By sending benign emails with the intent to lure the user into a false sense of security, threat actors lay the groundwork for a relationship to be more easily exploitable.

Proofpoint researchers observe multiple threat types sending benign emails to kickstart a conversation.

Lure and Task Business Email Compromise (BEC) threats typically start with a benign conversation or ask a question to get the recipient to engage with the email. Lure/task emails are typically a gateway theme – if the victim replies, they may be led to another type of threat such as a gift card, payroll, or invoice fraud. Proofpoint automatically identifies and blocks around 80,000 task themed emails each month.

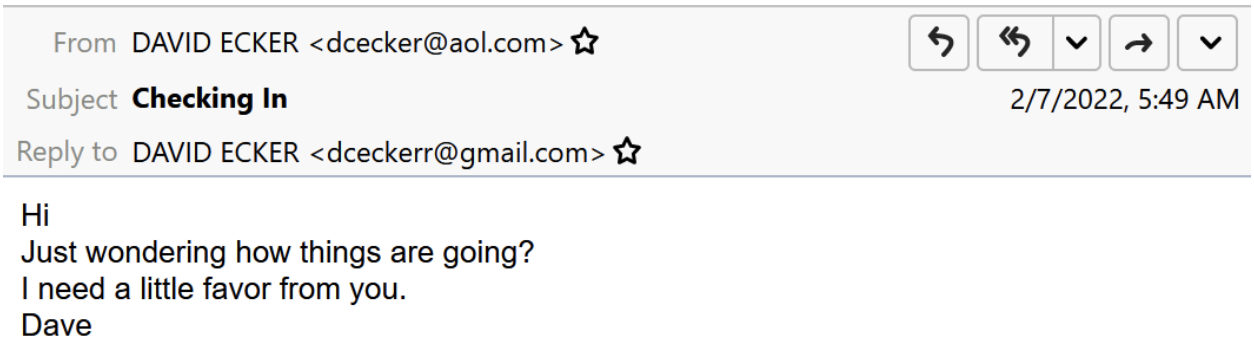

Figure: Lure/Task BEC theme email.

The threat actor tries to get a recipient to engage with them and will send follow up requests – such as transfering money – in future emails once a connection has been established. The result can cost individuals and organizations thousands of dollars.

Proofpoint researchers have also observed cybercriminal threat actors that distribute malware begin their interactions with victims using benign conversations. For example, the low-volume cybercriminal TA576 uses tax-themed lures specifically targeting accounting and finance organizations, and will email requests for tax preparation assistance. In campaigns observed in 2021, emails purported to come from "John Stevens" and his wife.

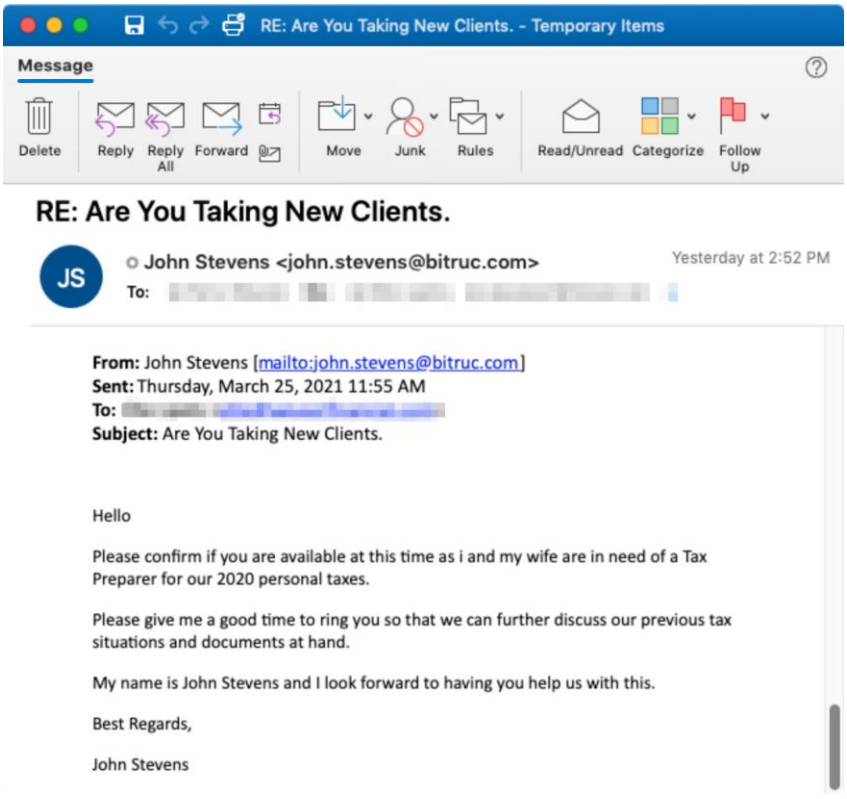

Figure: “John Stevens” requesting tax preparation assistance.

If the recipient replies, they receive a follow up email with a URL linking to a document that uses macros to drop a downloader that pulls in NetWire remote access trojan (RAT).

In 2021, ransomware threat actors used this method to engage with employees at large organizations to attempt to get them to install ransomware on their work computers and receive a cut of the ransomware profit. Proofpoint observed DEMONWARE ransomware threat actors using this method in the summer of 2021.

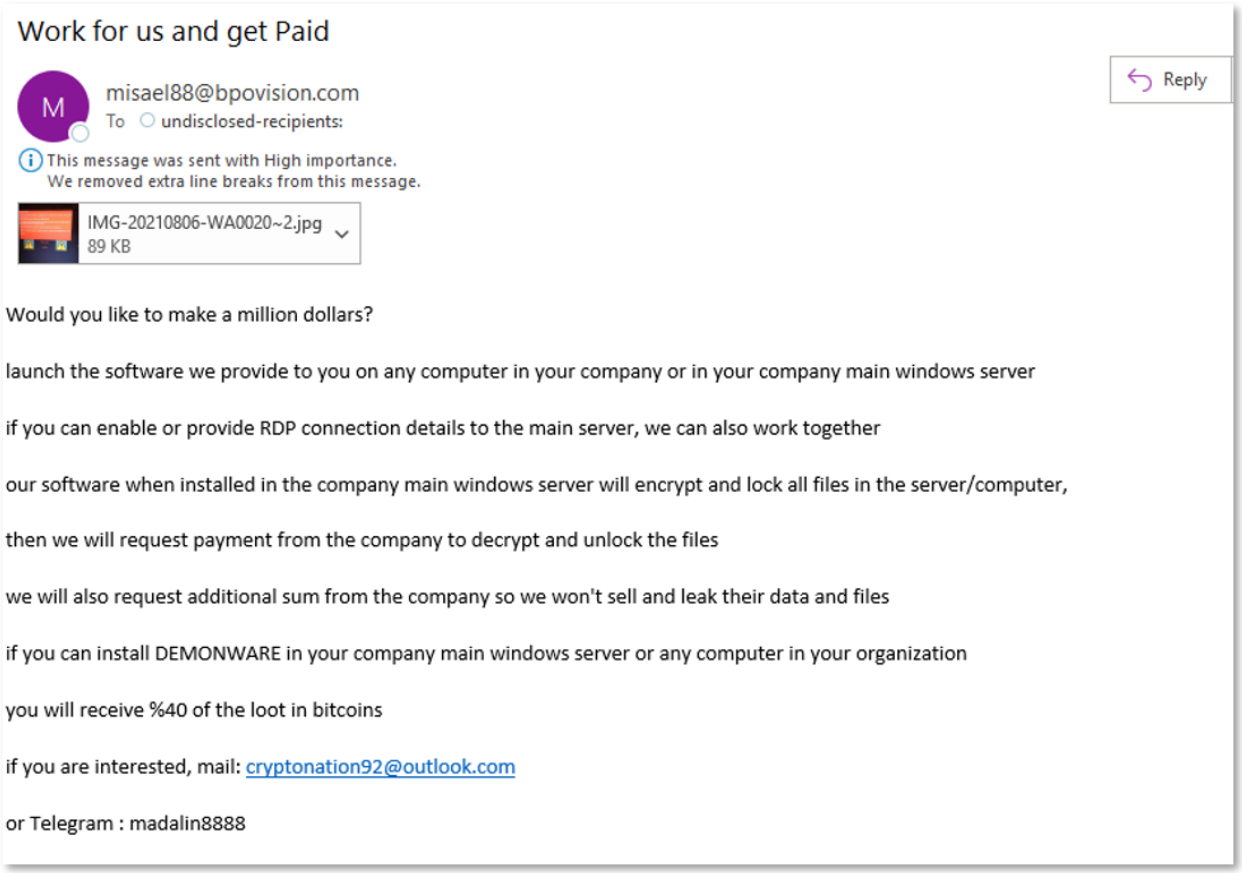

Figure: DEMONWARE ransomware extortion collaboration request.

The threat actor emailed recipients and requested the recipient enable RDP and/or collaborate directly to launch a ransomware attack. The threat actor offered 40% of ransom payment in bitcoin to the insider. The email did not contain malicious artifacts and instead directed them to a separate email and Telegram chat for further communications.

Advanced persistent threat (APT) actors that operate on behalf of state government interests are increasingly using benign conversations for relationship development to lay the groundwork for future attacks.

- TA453 campaigns have frequently used benign conversation, focused on building rapport to eventually solicit information from a target, specifically their login credentials in order to exfiltrate their mailbox. TA453, which is an Iran-aligned threat actor, has at times attempted to convince the target to share their cell phone number to communicate outside of email or their personal email address.
- In early 2021, TA406 began almost weekly campaigns featuring themes that included nuclear weapon safety, U.S. President

Joe Biden, Korean foreign policy and other political themes. TA406 is associated with the Democratic People's Republic of Korea (DPRK). The group attempted to collect credentials, such as Microsoft logins or other corporate credentials, from the targeted individuals. In some cases, the emails were benign in nature; these messages may have been attempts by the threat actor to engage with victims before sending them a malicious link or attachment.

- TA499, a Russia-aligned threat actor publicly known as Vovan and Lexus, began sending seemingly harmless, rapport building emails that attempt to solicit information from high-profile individuals in early 2021. The emails attempted to entice recipients into further contact via phone calls or remote video likely to create negative political content about the Russian opposition and policies in defiance of Russian President Vladimir Putin's objectives. TA499's emails masquerade as the team or wife of Alexei Navalny—a well-known Russian opposition leader—or Ukrainian officials.

Key Assumption: Content using legitimate services is safe

Users may be more inclined to interact with content if it appears to originate from a source they recognize and trust. However, threat actors regularly abuse legitimate services such as cloud storage providers and content distribution networks to host and distribute malware as well as credential harvesting portals. Overall, Google-related URLs were the most frequently abused in 2021 based on data from Proofpoint's Targeted Attack Protection (TAP) product.

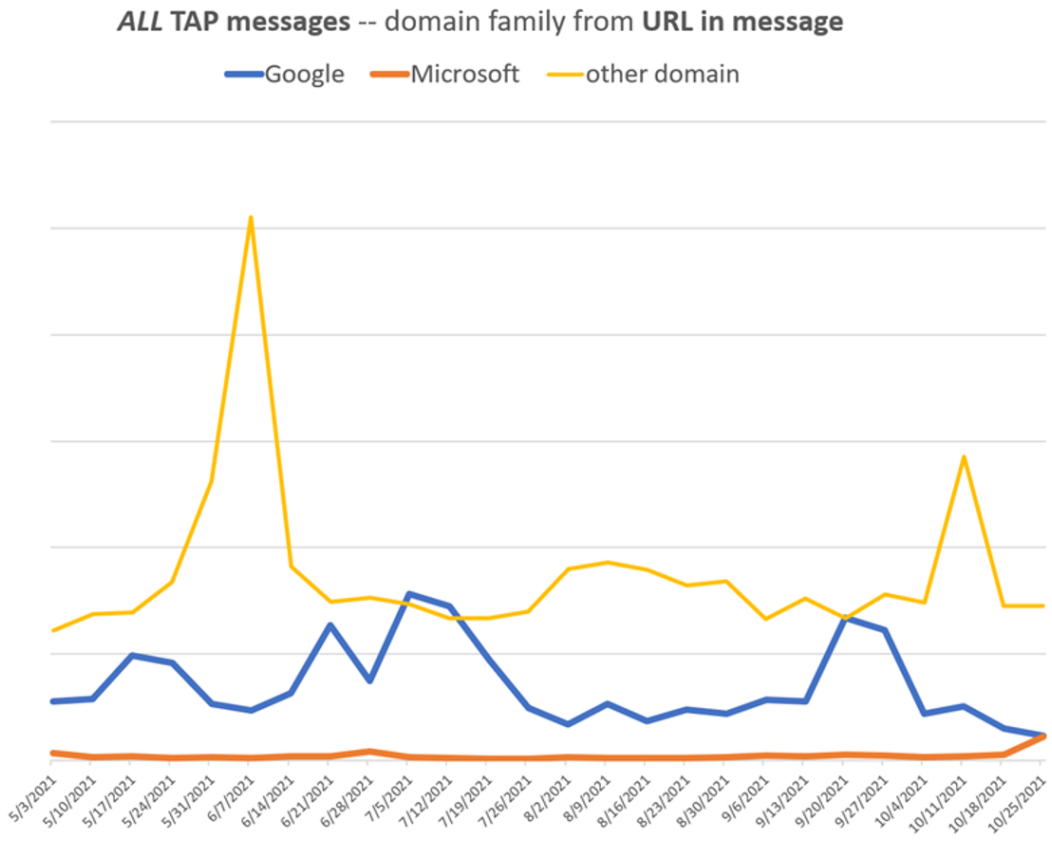

Figure: URL-based threats across all messages analyzed with TAP.

Despite this fact, when analyzing which domains are actually clicked, Microsoft-related URL-based threats earned more than twice the clicks of those hosted by Google. This could help explain why analysis at the campaign-level shows OneDrive is the most frequently abused service by top-tier e-crime actors, followed by Google Drive, Dropbox, Discord, Firebase, and SendGrid.

Proofpoint tracked nearly 1,000 campaigns leveraging these services in 2021. That accounts for approximately 14% of all observed campaigns.

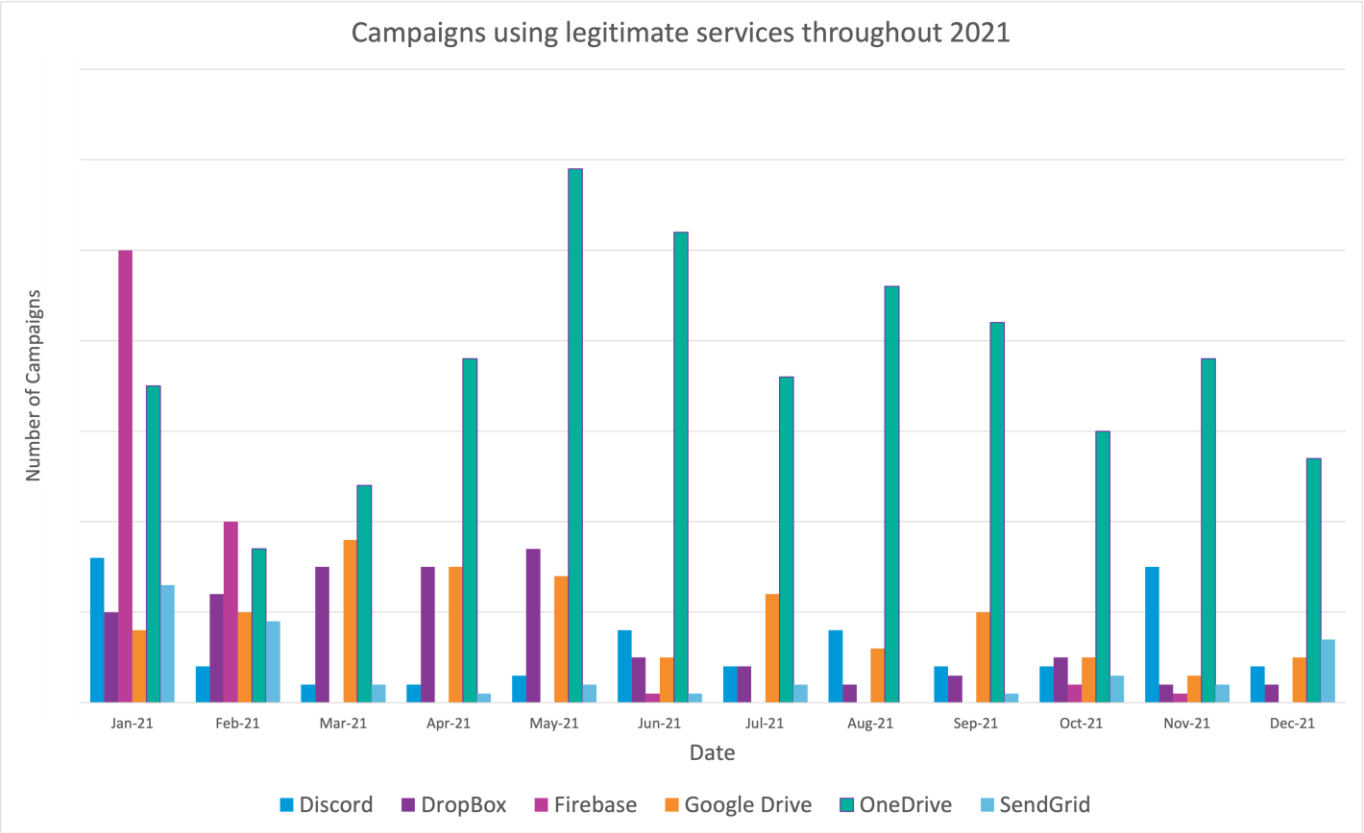

Figure: Campaigns using legitimate services in 2021.

Notably in 2021, threat actors consistently used Google's application development platform Firebase in campaigns throughout January and February, but it all but disappeared from identified campaigns during the rest of the year. Most of these campaigns were associated with credential phishing. It was the only legitimate service to be used more than OneDrive in January and February.

Attack paths leveraging remote services vary, but typically Proofpoint will see messages with malicious URLs directly in the message body or embedded in an attachment such as a PDF.

TA2541 for instance often sends emails containing Google Drive URLs that lead to an obfuscated Visual Basic Script (VBS) file. If executed, PowerShell pulls an executable from a text file hosted on various platforms such as Pastetext, Sharetext, and GitHub. The threat actor executes PowerShell into various Windows processes and queries Windows Management Instrumentation (WMI) for security products such as antivirus and firewall software, and attempts to disable built-in security protections. The threat actor will collect system information before downloading the RAT on the host.

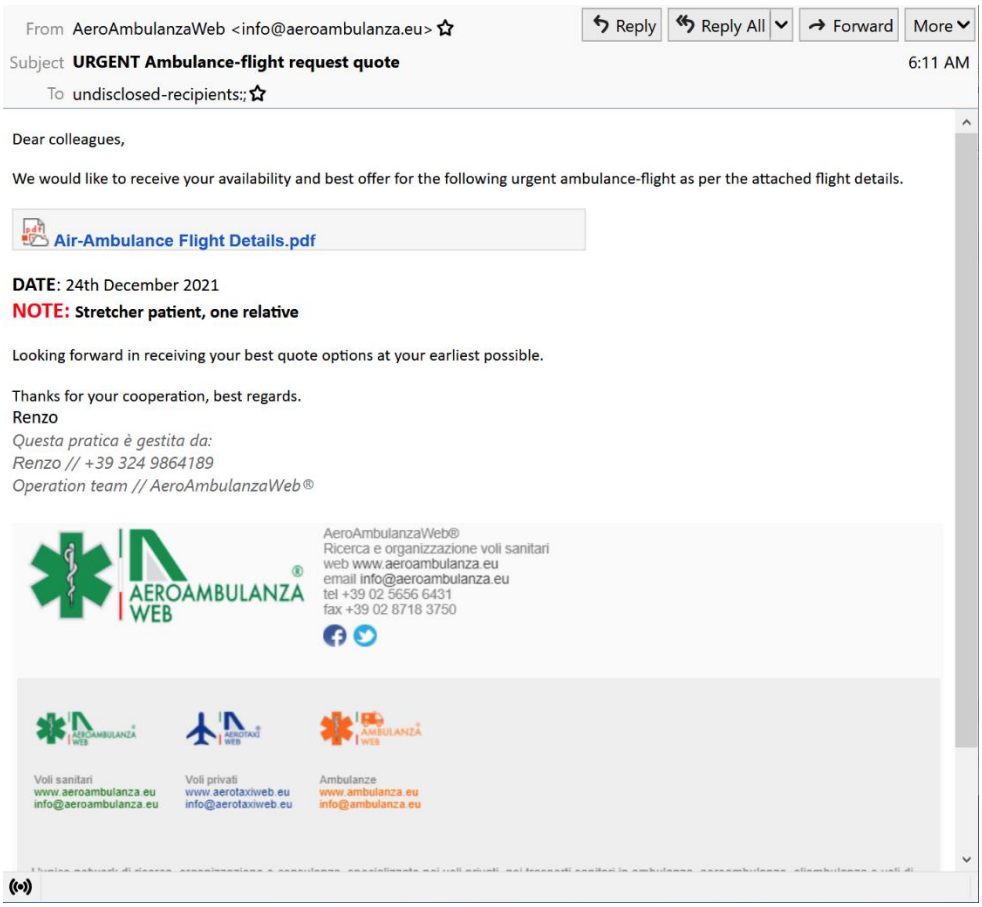

Figure: TA2541 lure.

Threat actors may prefer distributing malware via legitimate services due to their likelihood of bypassing security protections in email compared to malicious documents. Mitigating threats hosted on legitimate services continues to be a difficult vector to defend against as it likely involves implementation of a robust detection stack or policy-based blocking of services which might be business-relevant.

Key Assumption: Threat actors don't use the telephone

It's not unusual for people to think email-based threats live only in computers. But in 2021 Proofpoint researchers identified an increase in attacks perpetuated by threat actors leveraging a robust ecosystem of call center-based email threats. The threats are unique in that they require a lot of human interaction. The emails themselves don't contain malicious links or attachments, and individuals must proactively call a fake customer service number in the email to engage with the threat actor.

Proofpoint observes over 250,000 of these threat types each day.

There are two types of call center threat activity regularly observed by Proofpoint. One uses free, legitimate remote assistance software to steal money. The second leverages the use of malware disguised as a document to compromise a computer and can lead to follow-on malware. The second attack type is frequently associated with BazaLoader malware and is often referred to as BazaCall. Both attack types are what Proofpoint considers telephone-oriented attack delivery (TOAD).

The lures and themes threat actors send range in effort, from clearly fraudulent to using legitimate-looking brands and documents.

For example, our researcher identified a financially motivated TOAD threat masquerading as a PayPal invoice from a U.S. weapons manufacturer.

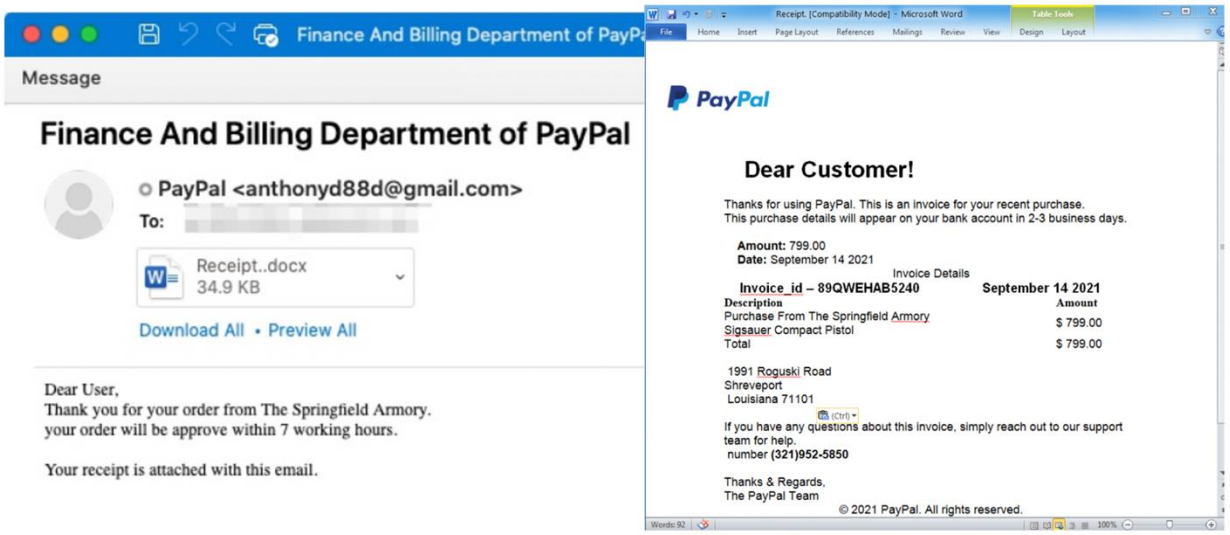

Figure: TOAD lure spoofing PayPal.

Victims can lose tens of thousands of dollars to these types of threats. In one case, Proofpoint is aware of a victim losing almost $50,000 to an attack from a threat actor purporting to be a Norton LifeLock representative.

Key Assumption: Replies to existing emails are safe

Thread hijacking, or conversation hijacking, is a technique where adversaries reply to existing benign email conversations with a malicious attachment, URL, or request to perform some action on the threat actor's behalf. An actor using this method preys on the person's trust in the existing email conversation. Typically, a recipient is expecting a reply from the sender, and is therefore more inclined to interact with the injected content.

To successfully hijack an existing conversation, threat actors need to obtain access to legitimate users' inboxes. This can be obtained in various ways including phishing, malware attacks, credential lists available on hacking forums, or password spraying techniques. Threat actors can also hijack entire email servers or mailboxes and automatically send replies from threat actor-controlled botnets.

In 2021, Proofpoint observed over 500 campaigns use thread hijacking, associated with 16 different malware families. Major threat actors including TA571, TA577, TA575 and TA542 regularly use thread hijacking in campaigns. In most observed cases, especially with threat actors distributing Qbot, Emotet, IcedID, and Ursnif, the thread hijacking mechanism is automated, in which email conversations are stolen off infected hosts and message replies are sent automatically by the attacker.

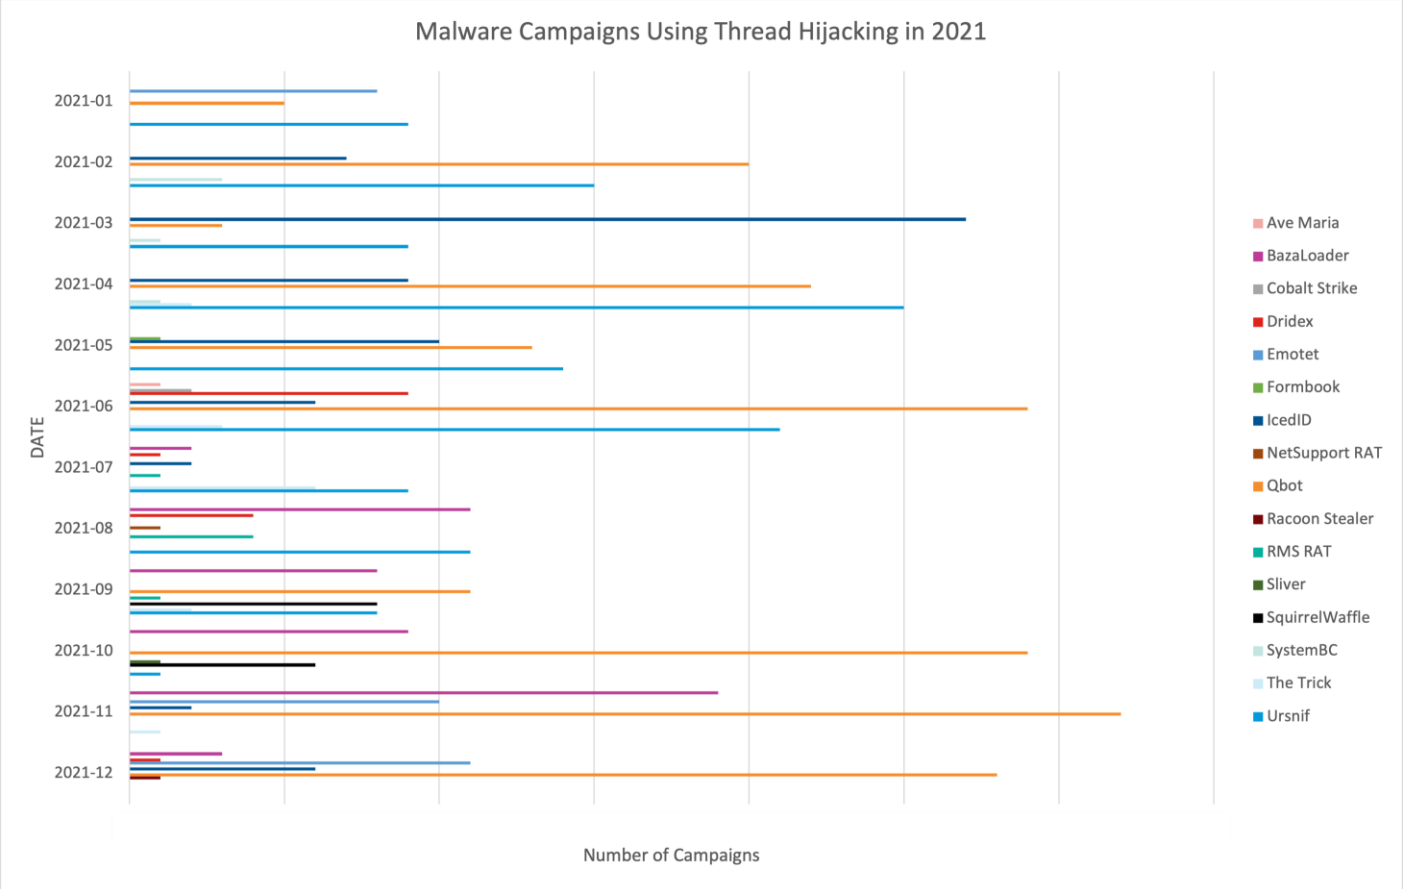

Figure: Malware campaigns using Thread Hijacking in 2021.

According to Proofpoint data, threat actors distributing high volume banking trojan campaigns consistently leveraged thread hijacking techniques more often than other threat actors. Specifically, threat actors distributing Qbot and Ursnif regularly rely on thread hijacking in campaigns. Many threat actors and associated trojans that leverage thread hijacking including IcedID, Qbot, or Dridex have been linked to follow-on malware campaigns including ransomware.

Typically email messages will look legitimate and because the threat is a reply to a legitimate thread, the message history will be attached. For example, in the following TA575 campaign, the threat actor replies to an existing thread purporting to provide payment details for the Administrative Office of the Illinois Courts (AOIC). The message contains a OneDrive link to a malware payload.

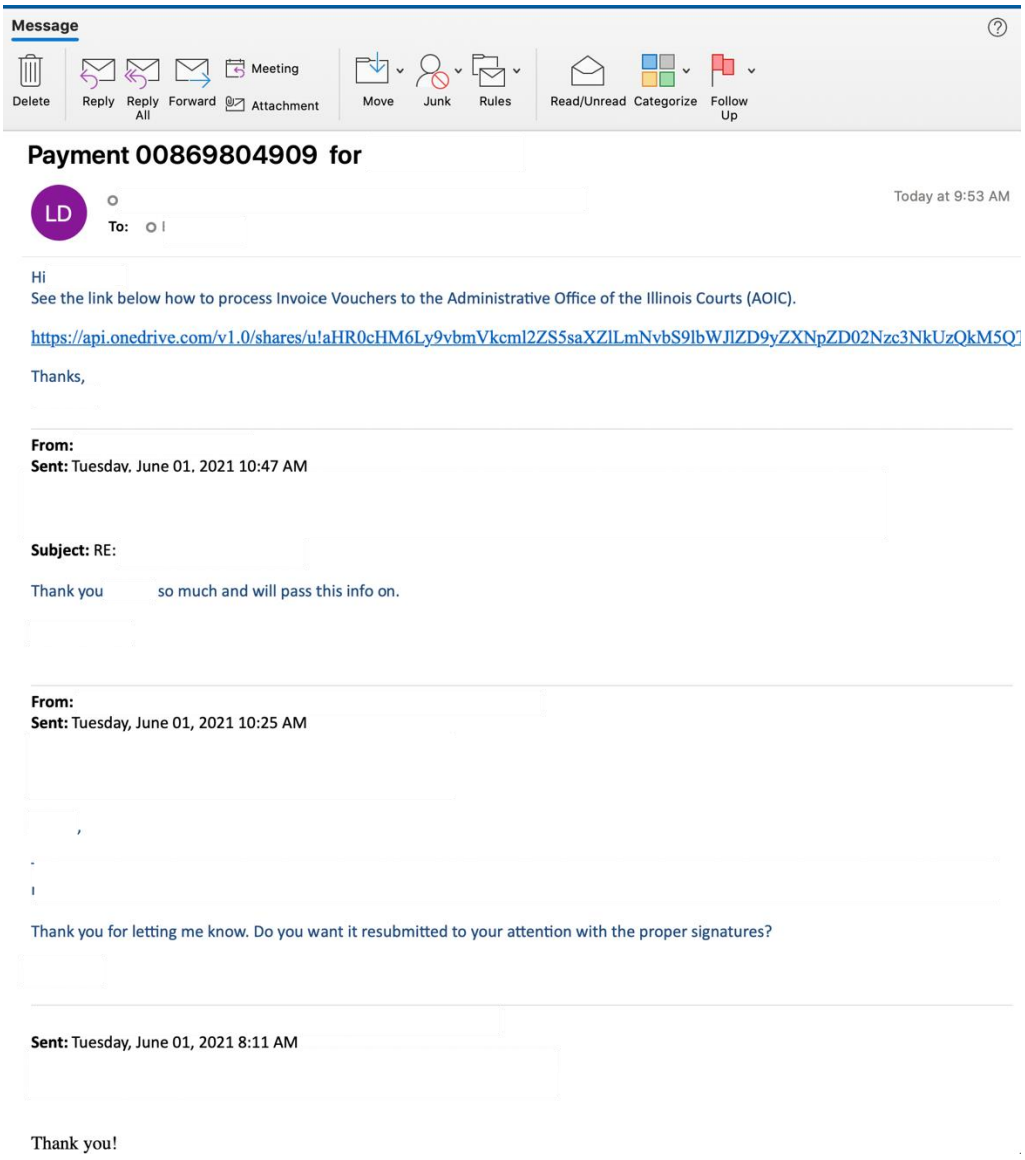

Figure: TA575 email lure.

The URL leads to a VBA macro-enabled Microsoft Word document spoofing AOIC. If enabled, the macros will download and execute Dridex malware.

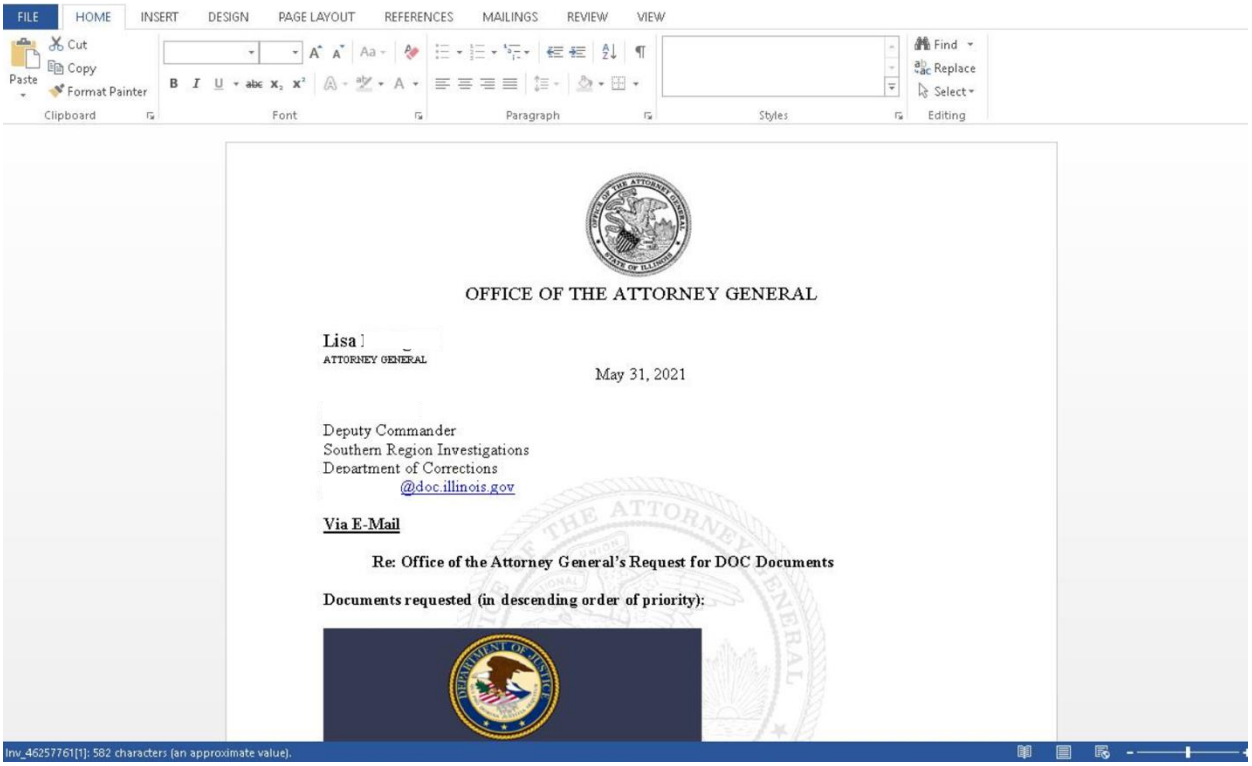

Figure: TA575 lure.

Business email compromise (BEC) threat actors also leverage thread hijacking, largely with invoice and payment themed lures. While they do not operate at the same scale as the banking trojans which have automated the process, the hands-on approach of tailoring a custom reply allows for a degree of personalization that is perhaps even more compelling.

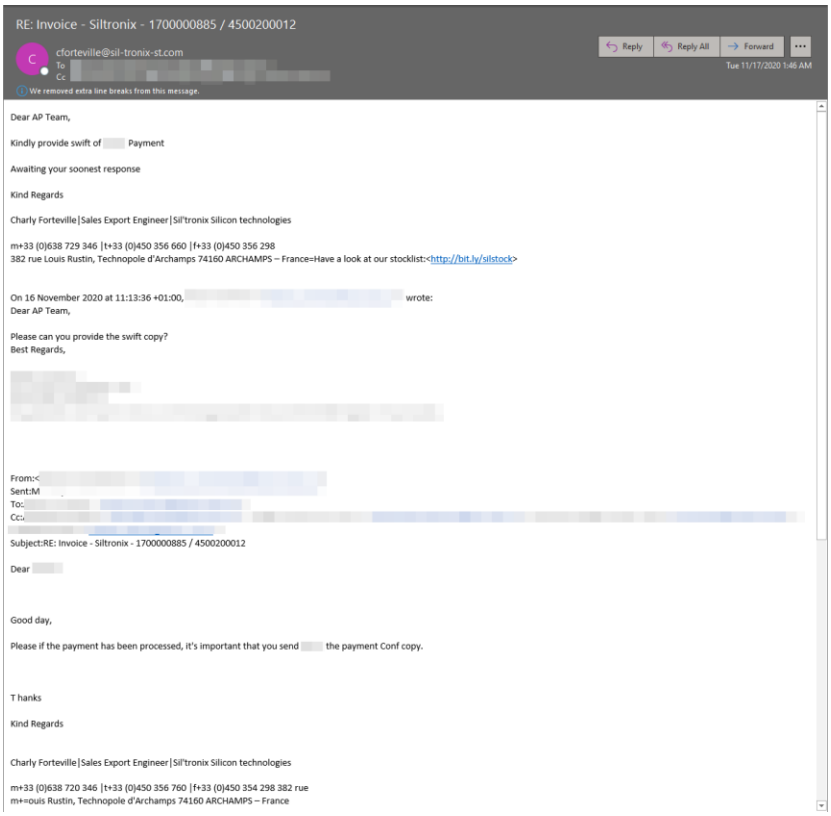
RE: Invoice - Siltronix - 1700000885 / 4500200012

cforteville@sil-tronix-st.com

Reply Reply All Forward

Tue 11/17/2020 1:46 AM

We removed extra line breaks from this message.

Dear AP Team,

Kindly provide swift of Payment

Awaiting your soonest response

Kind Regards

Charly Forteville|Sales Export Engineer|Sil'tronix Silicon technologies

m+33 (0)638 729 346 |t+33 (0)450 356 660 |f+33 (0)450 356 298
382 rue Louis Rustin, Technopole d'Archamps 74160 ARCHAMPS – France=Have a look at our stocklist:<http://bit.ly/silstock>

On 16 November 2020 at 11:13:36 +01:00, wrote:
Dear AP Team,

Please can you provide the swift copy?
Best Regards,

From:<
Sent:M
To:
Cc:
Subject:RE: Invoice - Siltronix - 1700000885 / 4500200012

Dear

Good day,

Please if the payment has been processed, it's important that you send the payment Conf copy.

T hanks

Kind Regards

Charly Forteville|Sales Export Engineer|Sil'tronix Silicon technologies

m+33 (0)638 720 346 |t+33 (0)450 356 760 |f+33 (0)450 354 298 382 rue
m+ouis Rustin, Technopole d'Archamps 74160 ARCHAMPS – France

Figure: BEC threat actor requesting payment in an existing thread.

Like most successful social engineering, this tactic relies on the trust users have in the authenticity of an email – in this case, that it is coming from a known good source. Automated processes make use of “RE:” and “FW:” in subject lines of stolen emails and then inject their own, generally unrelated, content into the thread. BEC threat actors leveraging thread hijacking will often respond to specific aspects of the discussion within the thread, increasing the likelihood that the recipient identifies the activity as legitimate. This circles back to our first key assumption: threat actors will hold in-depth conversations if they think it leads to a pay day.

Key Assumption: Threat actors only use business-related content

Every year threat actors capitalize on current events, news, and popular culture, using lure themes coinciding with things lots of people will be interested in to get people to engage with malicious content.

In January 2021, Proofpoint researchers observed a few BazaLoader campaigns leveraging Valentine's Day themes such as flowers and lingerie.

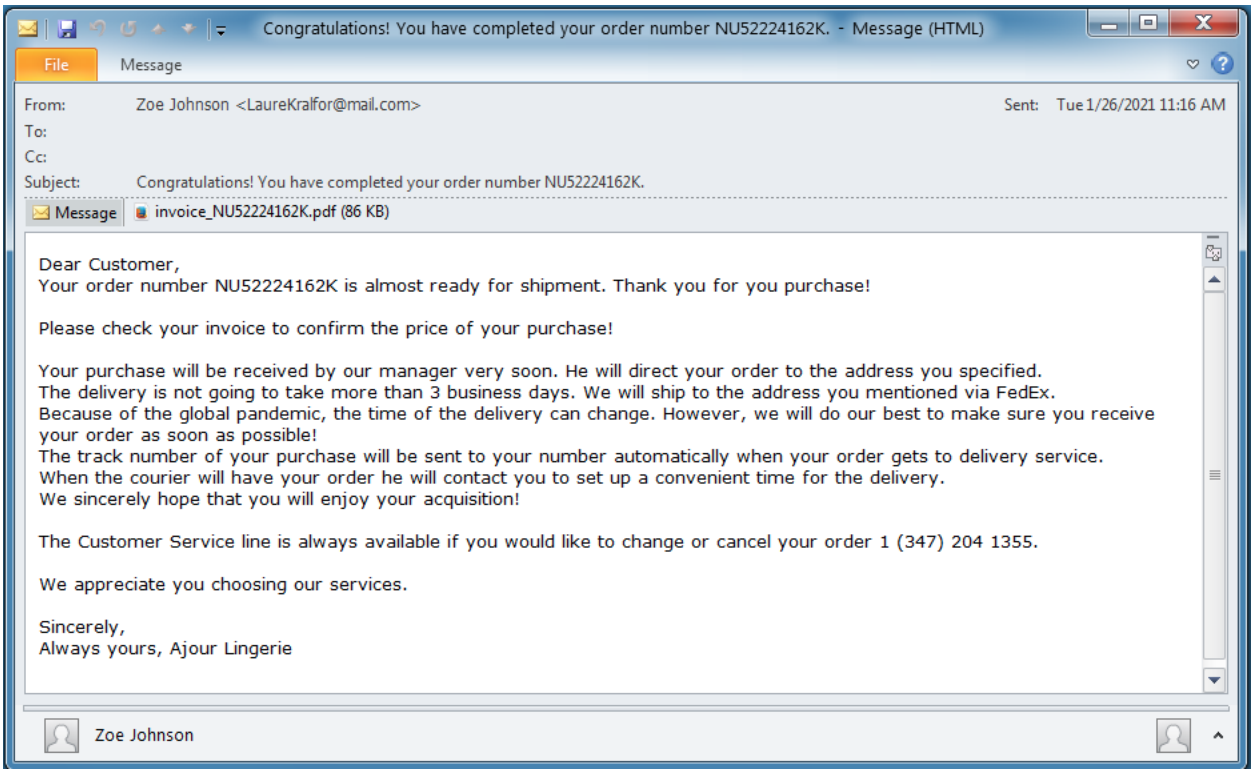

Figure: BazaLoader theme lure.

In 2021, BazaLoader threat actors began using infection chains that required a lot of human interaction, such as visiting actor-controlled websites to download the payload, or even calling the threat actor directly to presumably get assistance with an erroneous purchase.

Dridex malware distributors used pop culture themed lures in 2021, too, attempting to capitalize on the popularity of the Netflix smash hit Squid Game. In October 2021, Proofpoint identified the large cybercrime actor TA575 distributing the Dridex banking trojan using Squid Game themes targeting users in the U.S. The threat actor purported to be entities associated with the Netflix global phenomenon using emails enticing targets to get early access to a new season of Squid Game or to become a part of the TV show casting.

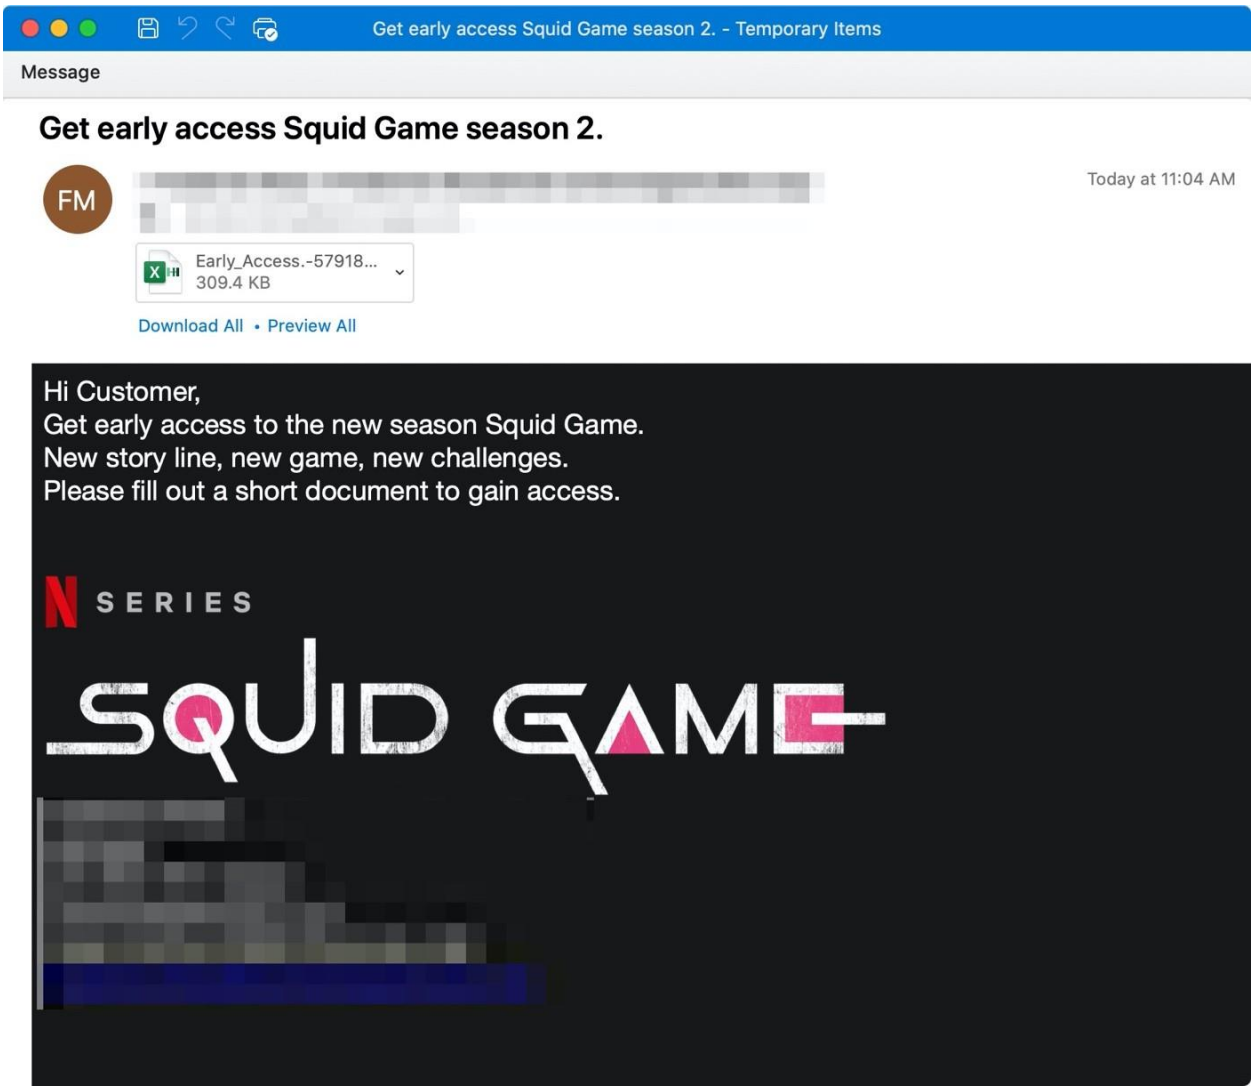

Figure: Squid Game lures.

Tax themes are a regular favorite of cybercriminal threat actors.

One of the most convincing Internal Revenue Service (IRS)-themed campaigns leveraged the idea that the potential victim was owed an additional refund in an attempt to harvest a variety of personally identifying information (PII). This included pieces like the previous years' adjusted gross income and PIN which could allow the threat actors to attempt to claim the victims' refunds.

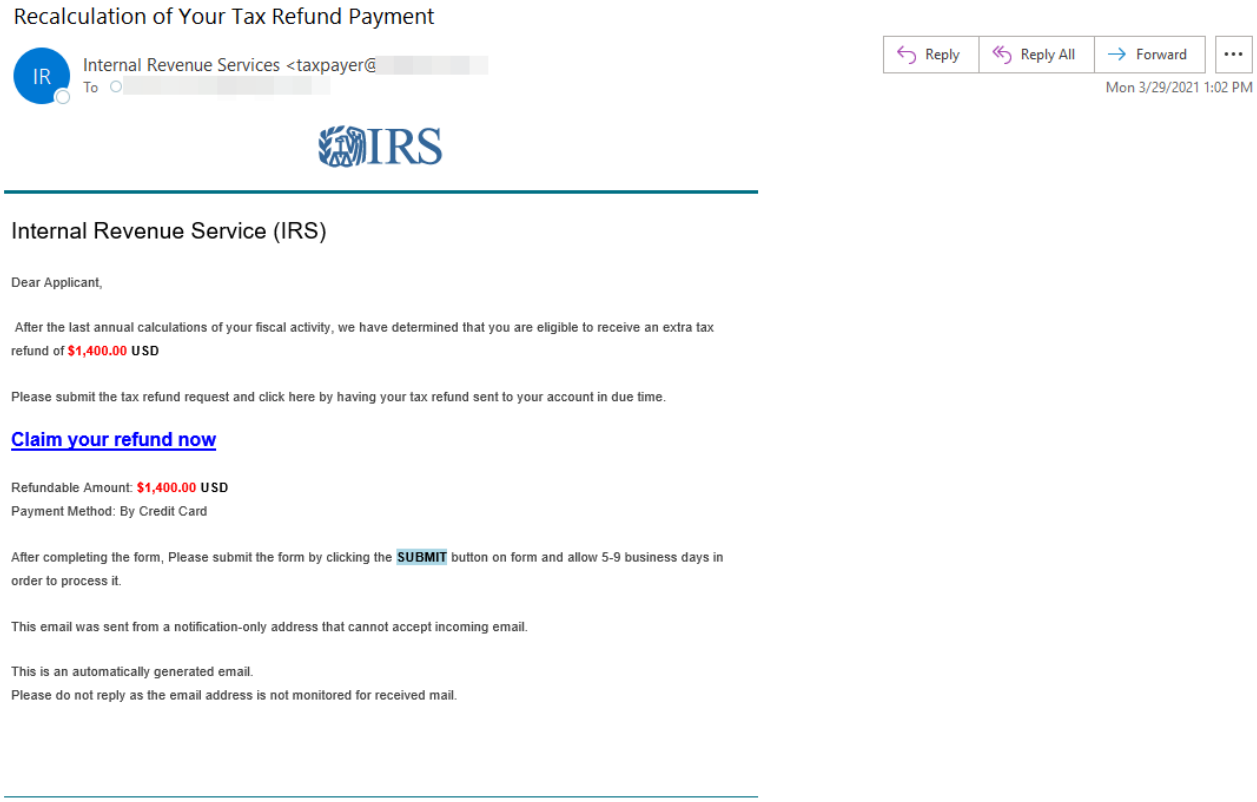

Recalculation of Your Tax Refund Payment

IR Internal Revenue Services <taxpayer@
To

Reply | Reply All | Forward | ...
Mon 3/29/2021 1:02 PM

IRS

Internal Revenue Service (IRS)

Dear Applicant,

After the last annual calculations of your fiscal activity, we have determined that you are eligible to receive an extra tax refund of $1,400.00 USD

Please submit the tax refund request and click here by having your tax refund sent to your account in due time.

Claim your refund now

Refundable Amount: $1,400.00 USD
Payment Method: By Credit Card

After completing the form, Please submit the form by clicking the SUBMIT button on form and allow 5-9 business days in order to process it.

This email was sent from a notification-only address that cannot accept incoming email.

This is an automatically generated email.
Please do not reply as the email address is not monitored for received mail.

Figure: IRS tax theme lure.

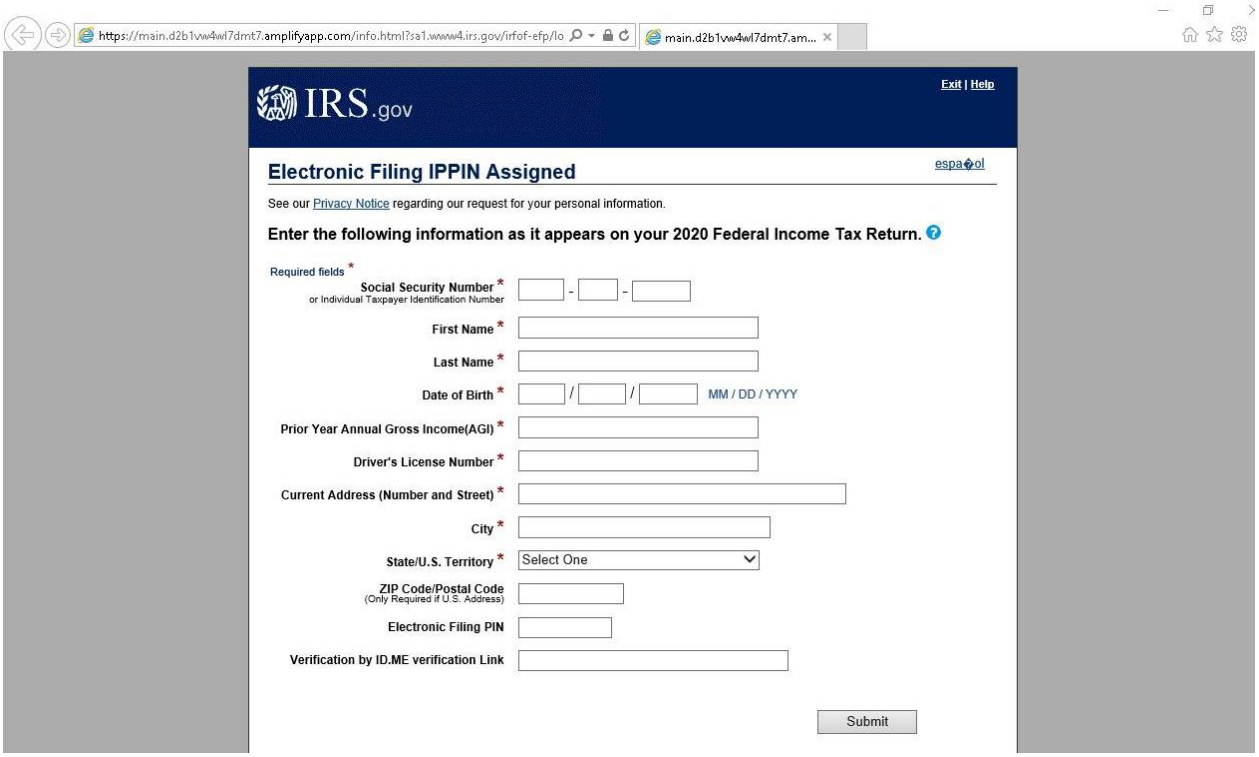

Figure: IRS tax theme landing page.

COVID-19

2021 saw the second full year of a global pandemic, and although it remained top of mind for most people, threat actors used pandemic themes considerably less compared to 2020. Threat actor use of COVID-19 themes in 2021 appeared mainly driven by mainstream narratives around vaccines, corporate responses to vaccines and masks, and the Delta and Omicron variants. Comparatively, 2020 actor abuse of COVID-19 content was more generally reactive in general to the burgeoning narrative.

COVID-19 remains a very important topic, which people have strong feelings about. Such topics can be good lure themes for threat actors as people are more likely to engage with content on topics they feel strongly about. However, it's likely due to public reporting and user education, people became more aware of threat actors using COVID-19 themes for malicious activities, and thus they became less effective.

On average, Proofpoint observed over six million COVID-19 related threats per day through 2021. The data demonstrate that the topic was both omnipresent and also slightly declined in usage over the course of the year.

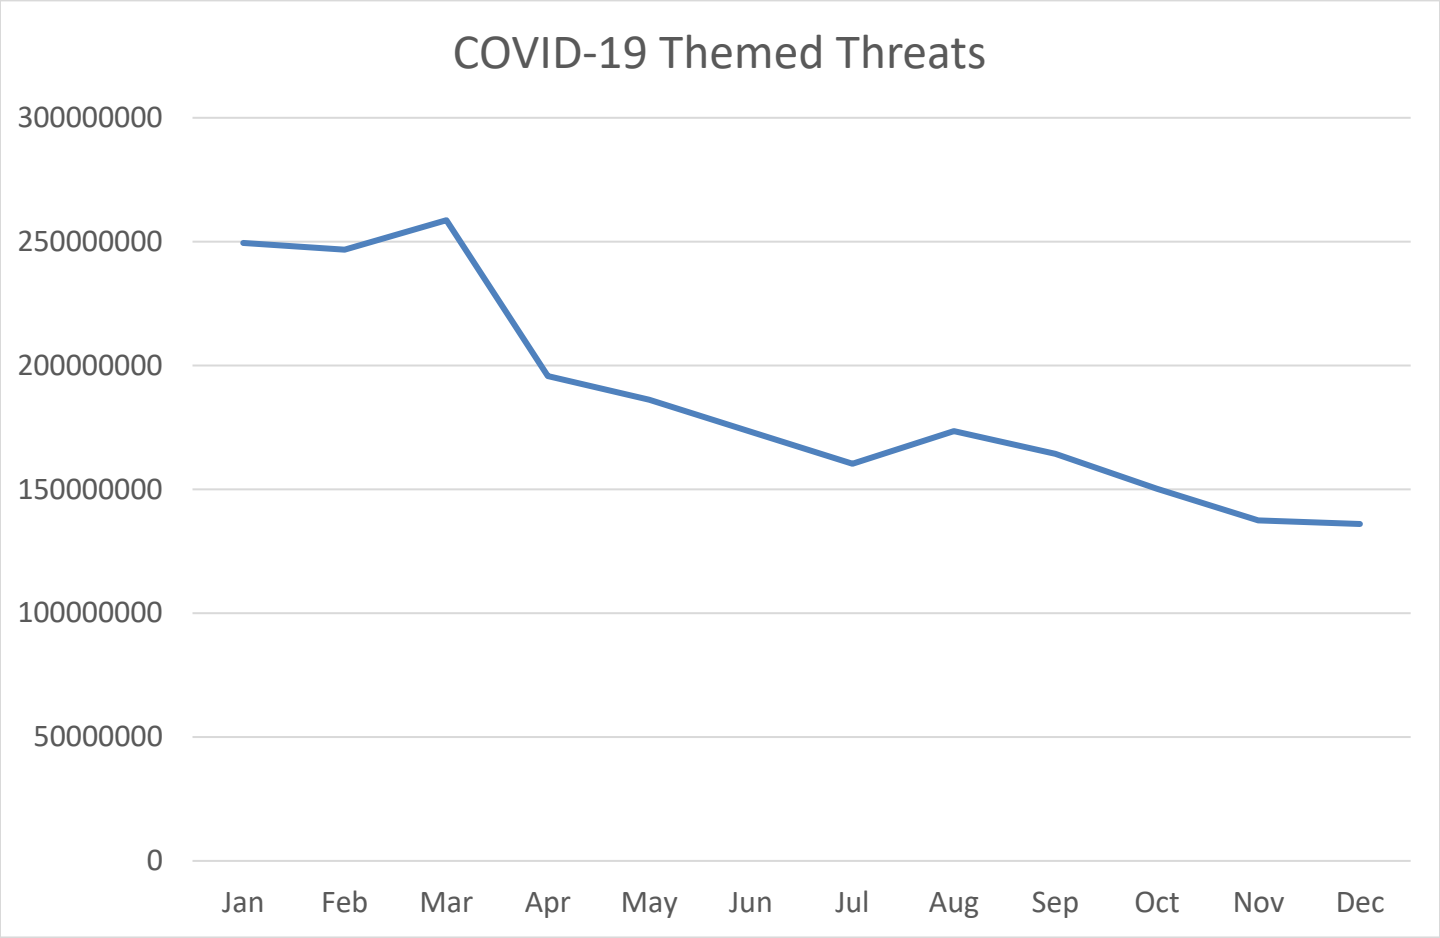

Figure: Amount of COVID-19 related threats observed monthly

Proofpoint observed cybercrime, BEC, and APT threat actors using COVID-19 themed lures.

January 2021 kicked off with a high-volume campaign spoofing the World Health Organization (WHO), an organization commonly used by threat actors associated with COVID-19 themes. The messages contained a URL which led to a fake WHO authentication page designed to harvest user credentials. After POST action the user was redirected to a login page on careers[.]who[.]int.

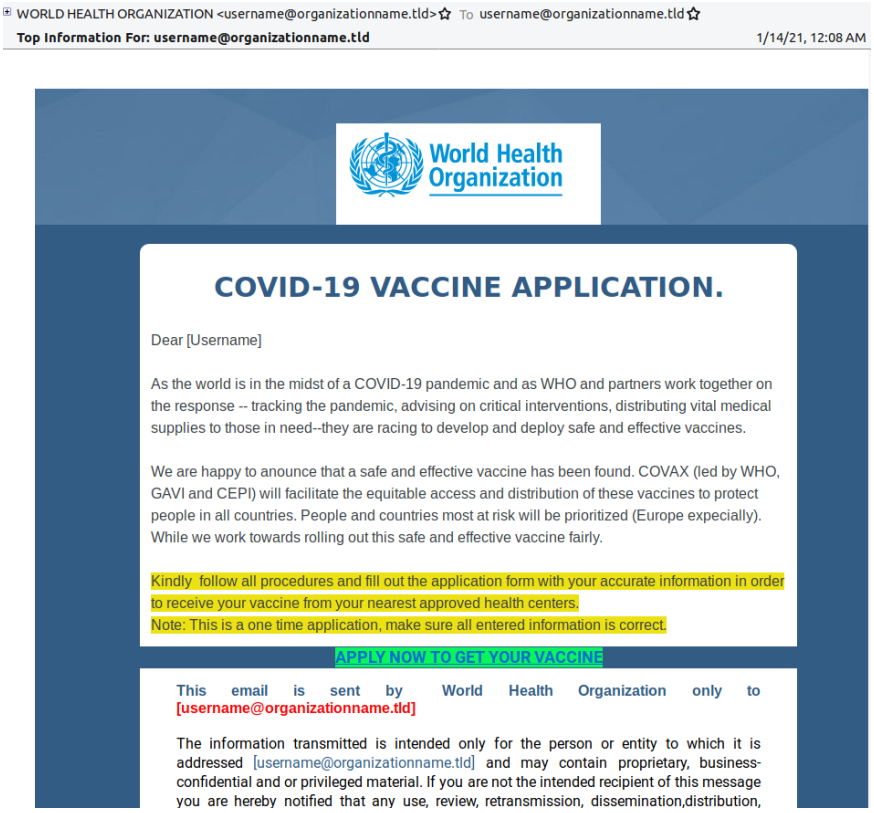

Figure: WHO-theme lure.

In September 2021, Proofpoint identified TA3546, also known as FIN7, leveraging COVID-19 themes to distribute GRIFFON malware.

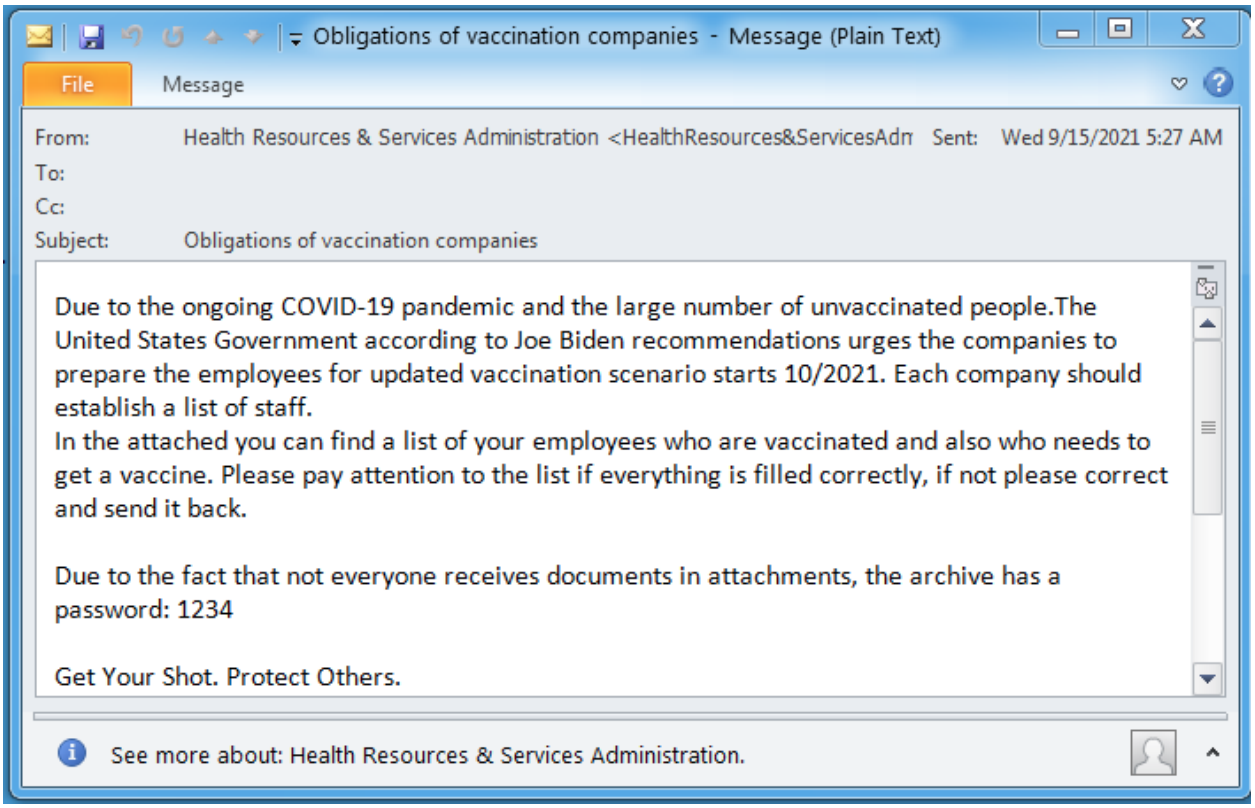

Figure: TA3546 lure.

The messages purported to be from a Health Resources & Services Administration. The messages contained zipped JavaScript attachments that installed the GRIFFON backdoor, which downloads additional code for system profiling.

From autumn through late 2021, Proofpoint researchers identified an increase in email threats targeting mostly North American universities attempting to steal university login credentials. The threats typically leverage COVID-19 themes including testing information and the new Omicron variant.

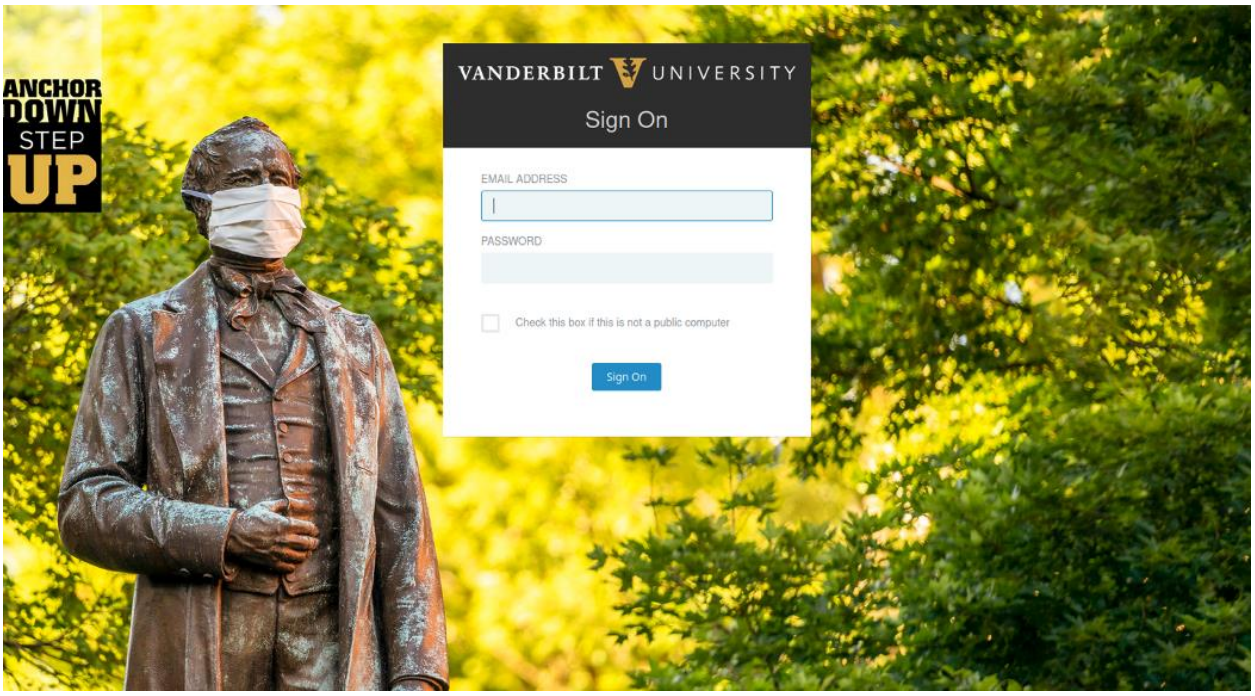

Figure: Credential capture portal spoofing a university web page.

The phishing emails contained attachments or URLs for pages intended to harvest credentials for university accounts. The landing pages typically imitated the university's official login portal, although some campaigns featured generic Office 365 login portals. Some lures used lures specific to the Omicron variant.

Early in 2021, Proofpoint researchers observed the Iran-aligned APT actor TA451 (APT33) using COVID-themed lures in a phishing campaign against a US defense contractor. Masquerading as the World Health Organization, the actor delivered malicious emails with a link to a COVID19tracker[.]exe file. The executable reached out to download a batch script (iehchecker[.]bat) that downloaded a PowerShell script (Update-KB4524147[.]ps1) with reverse shell capabilities.

By late 2021, Proofpoint researchers had observed several more APT actors using COVID lures in their campaigns. The Russia state-sponsored TA421, publicly known as APT29, targeted various government entities worldwide with COVID lures that delivered an HTML which constructed an ISO file that ultimately led to the delivery of Cobalt Strike. Meanwhile, the Iran-aligned TA456, also known as Tortoiseshell, used Omicron COVID-19 themed emails in reconnaissance and profiling campaigns targeting academics. And, in yet another campaign, an Indian APT actor, tracked by Proofpoint as TA425, distributed emails with a COVID-19 booster shot lure targeting users in Pakistan. The landing page in this campaign impersonated the Pakistani National Immunization Management System and hosted a password protected macro-laden Excel file which dropped xRAT—a legitimate remote administration tool.

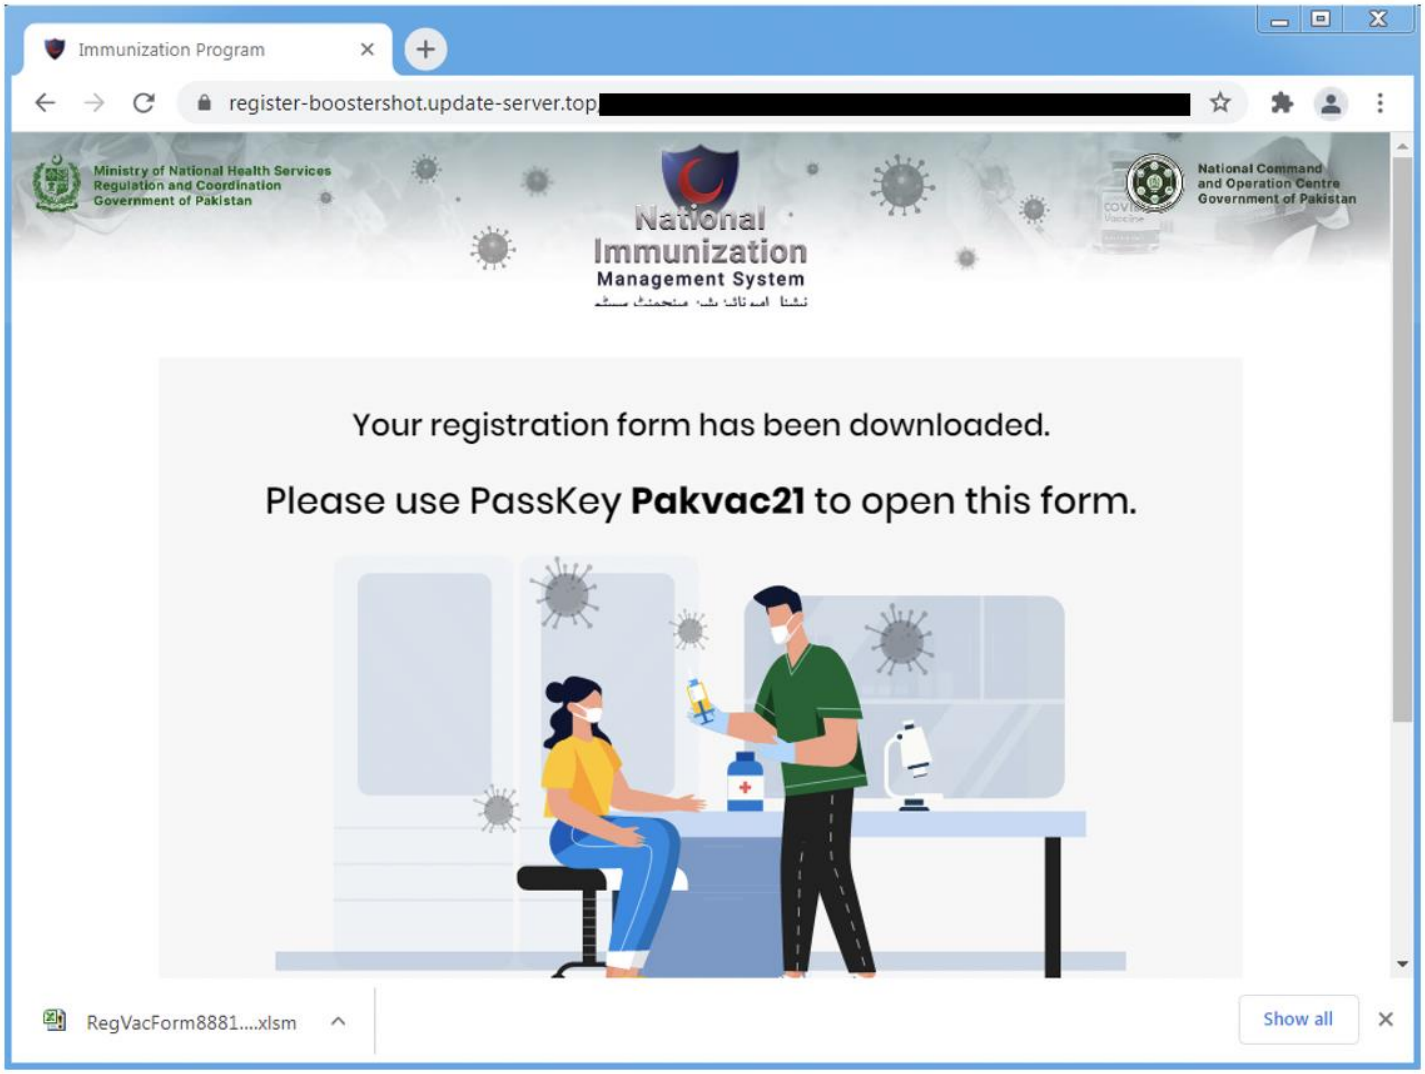

Figure: Landing page spoofing Pakistani National Immunization Management System.

Figure: Macro-laden Excel document used to download and execute xRAT.

Looking Forward

The driving force behind the widespread use of social engineering is the fact that it is effective -- despite defenders' best efforts, cybercriminals continue to be successful at exploiting the human element to recognize financial gain. This is unlikely to change any time soon. The most sophisticated criminal organizations have evolved to mirror legitimate businesses and as a result have scaled to become more resilient while also recognizing greater profits than ever before. Until some factor creates a situation where the path of least resistance to monetization is not a person, threat actors will continue to capitalize by preying on human behaviors, instincts, and emotions.

As outlined in this report, some notable ways threat actors achieved this in 2021 was by undermining behavior resulting from faulty assumptions on the part of the end-user:

- The assumption that threat actors won't hold in-depth conversations

- The assumption that legitimate services are always safe to use
- The assumption that threat actors won't use technologies beyond the computer (like the telephone)
- The assumption that threat actors are unaware of existing conversations between colleagues and won't leverage them
- The assumption that threat actors only use business-related content

Looking forward, there are two essential questions: "What will threat actors use next?" and "What can we do to make threat actors less successful?".

Threat actors have demonstrated time after time that they will use what works. Defenders can assume that they will continue to regularly see receipts, invoices, documents, and spreadsheets – files essential for the day-to-day operation of businesses. Likewise, every year Proofpoint researchers observe campaigns based on calendar-year cyclical events like holidays and tax season.

Newer developments demonstrate a point often lost in the race to stay ahead of the adversary but one that is also key to Proofpoint's people-centric security model: threat actors are people too. They are immersed in the same social and political narratives we all are. They know what is popular. They know what we care about. It's not surprising, as a result, to observe campaigns themed around trending television shows or protest movements. To answer the question "What will threat actors use next?", ask yourself: what is the thing we are all talking about right now?

Finally, what can we do about it? For many organizations, infrastructure has been sufficiently hardened. Policy and law enforcement actions continue to disrupt criminal organizations where possible, even though resilient groups often return to malicious activity in a matter of months. All evidence points to the end-user as the weakest point in a modern layered defense. Logically, defenders must address the weakest point in the system. To that end, many organizations have implemented security awareness programs. Too often, at the end-user level, these programs are treated in the same fashion as "compliance training": namely, completion is a box to be checked off and never thought of again. The most impactful course of action, for any given organization, is to shift the culture toward a posture where identification of incoming threats is understood as both relevant and necessary day-to-day. Likely, this means encouraging familiarization with the wide array of content threat actors may leverage and imposing few obstacles to more regular flagging of content as potentially malicious (and no repercussions for cleared content).

Organizations must ingrain in their users the idea that malicious activity is regular, even inevitable. As this becomes more widely accepted and reporting/clearing pipelines for threats become more well-established within workflows, threat actors should have a progressively more difficult task in exploiting the human element.

LEARN MORE

For more information, visit **proofpoint.com**.