CLOUDFLARE

2023 Phishing Threats Report

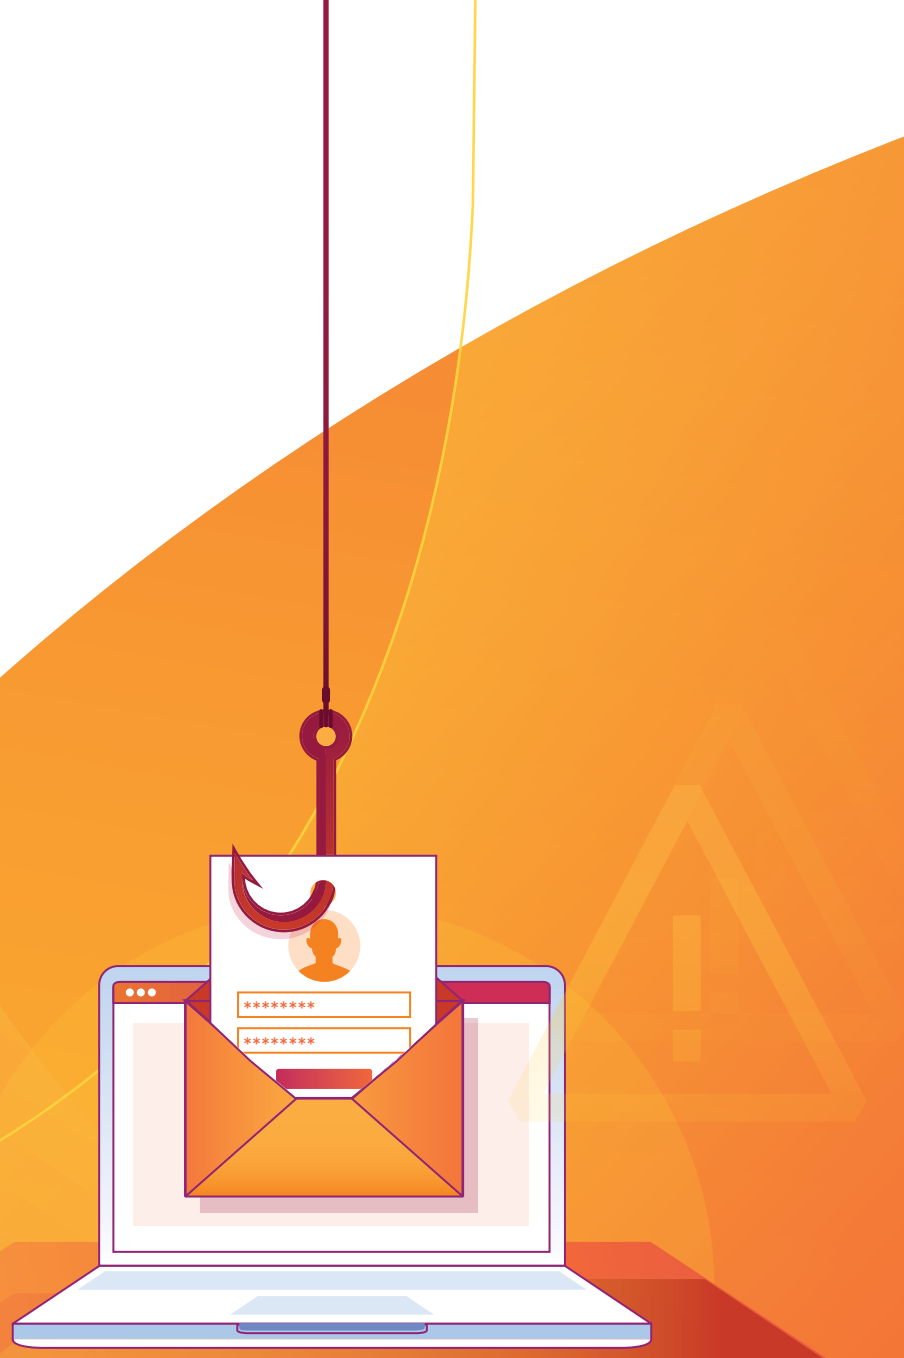

Table of Contents

Click section to skip to its page

About this report

Table of contents

Email is the most exploited business application. It is the primary initial attack vector for cybersecurity incidents, and contains vast amounts of trade secrets, PII, financial data, and other sensitive matters of value to attackers.

On top of that, email is one of the hardest applications to secure. If it were simple, there would be fewer headlines about business email compromise (BEC) losses topping $50 billion[1], and fewer breaches resulting from someone falling for a phish. Once an attacker has infiltrated one email account, they can move laterally and impact a wide range of internal systems.

To examine key phishing trends, this inaugural **Cloudflare Phishing Threats Report** is based on threat intelligence incorporating data from the 112 billion threats that Cloudflare's global network blocks daily. For this report's purpose, we evaluated a sample of more than **279 million email threat indicators[2], 250 million malicious messages[3], nearly 1 billion instances of brand impersonation[4], and other data points gathered from approximately 13 billion emails processed between May 2022 to May 2023**.

Additionally, this report is informed by a Cloudflare-commissioned study conducted by Forrester Consulting. Between January 2023 and February 2023, **Forrester Consulting surveyed 316 security decision-makers across North America, EMEA, and APAC[5]** about the state of phishing.

In the following pages we explore these three takeaways from our research:

- **Attackers use links as the #1 phishing tactic** — and are evolving how they get you to click and when they weaponize the link.
- **Identity deception** takes multiple forms **and can easily bypass email authentication standards**.
- Attackers may pretend to be hundreds of different organizations, but they **primarily impersonate the entities we trust (and need to get work done)**.

But make no mistake: **attackers don't just go after businesses**. For example, we observed more messages impersonating the United Nations than the New York Stock Exchange[4]. And in the three months leading up to the 2022 US midterm elections, we prevented around 150,000 phishing emails targeting campaign officials.

We hope our findings and recommendations help you tackle the key component underpinning phishing attacks: **trust**.

Trust that the person or entity you're communicating with is who they say they are, that what they are sharing is legitimate, and that their communication channel—how they contact you—has not been compromised.

Key findings

Table of contents

#1 method

Deceptive links were the #1 method for cyber actors, comprising 35.6% of threats.[2]

89%

Email authentication doesn't stop threats. The majority (89%) of unwanted messages "passed" SPF, DKIM, or DMARC checks[8].

Over 1,000 organizations

Attackers posed as more than 1,000 different organizations in their brand impersonation attempts. However, in the majority (51.7%) of incidents, they impersonated one of 20 of the largest global brands.[4]

#2 threat category

One-third (30%) of detected threats featured newly registered domains — the #2 threat category.[7]

39.6 million

Identity deception threats are on the rise — increasing YoY from 10.3% to 14.2% (39.6 million) of total threat indicators[6]

Trusted companies

The most impersonated brand happens to be one of the most trusted software companies: Microsoft. Other top companies impersonated included Google, Salesforce, Notion.so, and more.[4]

Multichannel phishing threats

90% of surveyed security decision-makers agree that the type and scope of phishing threats is expanding — with 89% concerned about multichannel phishing threats.[5]

Top email threats overview

Table of contents

Below is a snapshot of the top email threat categories we observed between May 2, 2022 - May 2, 2023.[2]

Attackers often use a combination of social engineering and technical tactics to make their messages seem legitimate to both the recipient and the recipient's email systems. Therefore, we look at many data signals to identify and block unwanted emails. These signals include:

- **Structural analysis** of headers, body copy, images, links, attachments, payloads, and more, using heuristics and machine learning models specifically designed for these signals
- **Sentiment analysis** to detect changes in patterns and behaviors (e.g., writing patterns and expressions)
- **Trust graphs** that evaluate partner social graphs, sending history, and potential partner impersonations

From there, we categorize threat indicators into **over 30** different types.

Read on for more insights about these key categories in particular: Deceptive links, domain age, identity deception, brand impersonation, account compromise, and BEC.

Detections by threat category

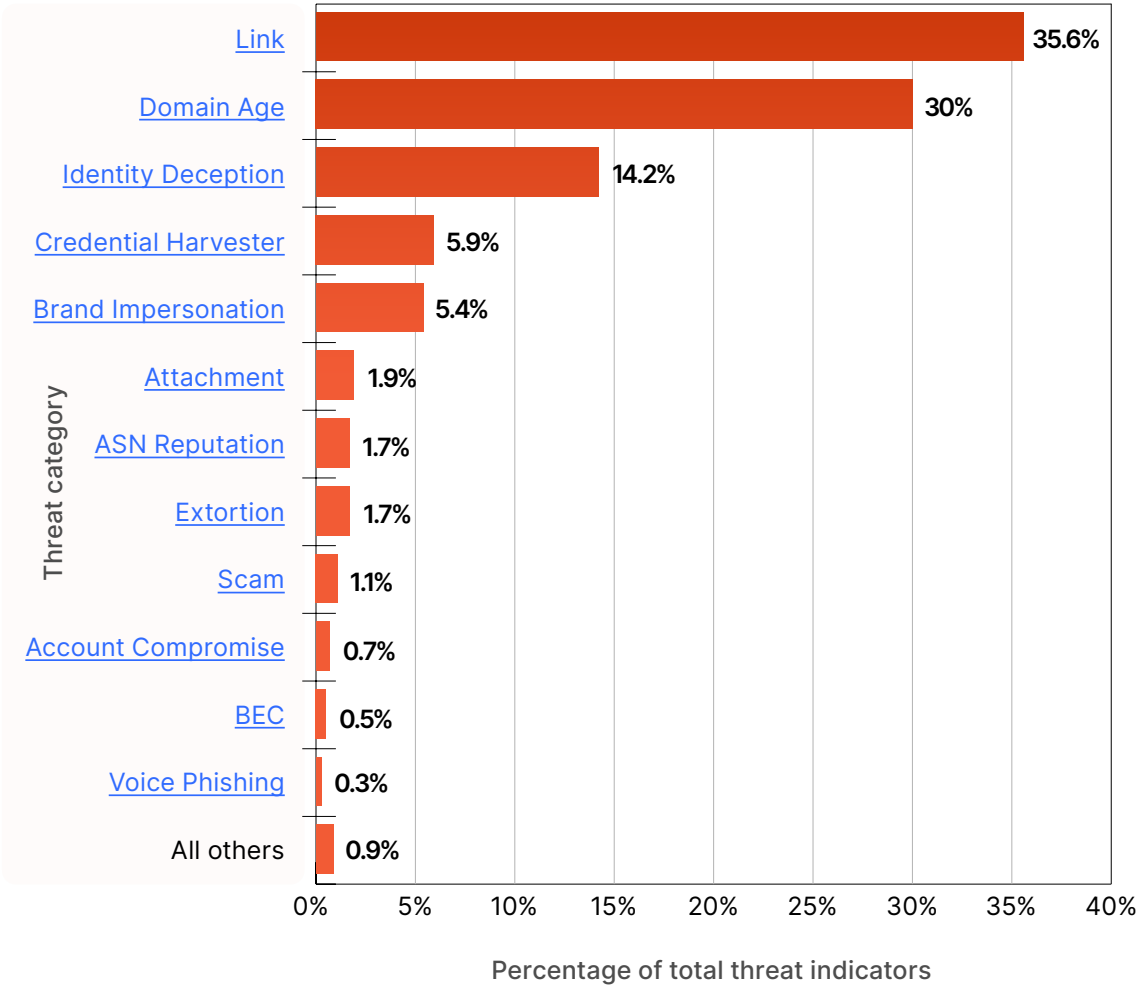

Detailed descriptions of the above-noted categories can be found in the Appendix.

Top Threat: Deceptive links

Table of contents

Deceptive links were the #1 email threat category — appearing in 35.6% of our detections. Links were also the #1 threat category the prior year (May 2021 - April 2022), when they comprised 38.4% of all threat indicators.

It's natural to want to interact with a link from someone you "know" — especially if it's timely and looks like prior emails. But clicking the wrong link can lead to consequences such as:

- **Credential harvesting** if you enter credentials on an attacker-controlled page
- **Remote code execution** (RCE) that lets the attacker install malware or ransomware, steal data, or take other actions
- **Network compromise** from taking over one workstation

People still click because it's in our nature. As the Verizon 2023 Data Breach Investigations Report (DBIR) notes,

"the human element still makes up the overwhelming majority of incidents, and is a factor in 74% of total breaches."[12]

Look closer: Scammers exploit a crisis in real time

This DocuSign-themed March 2023 SVB campaign, which targeted dozens of individuals at multiple organizations (including Cloudflare's co-founder and CEO, Matthew Prince), included HTML code that contains an initial link and a complex redirect chain that is four-deep.

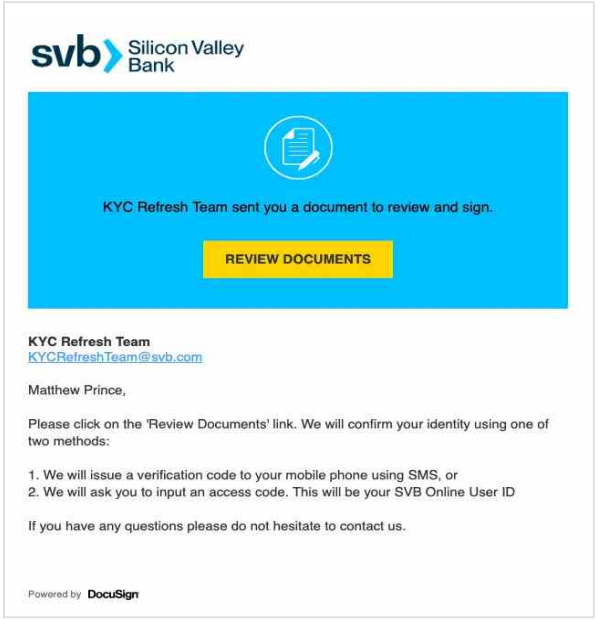

We automatically blocked this campaign for Cloudflare email security customers, but the chain would begin if a user clicks the 'Review Documents' link. It takes the user to a trackable analytic link run by Sizmek by Amazon Advertising Server bs[.]serving-sys[.]com.

The link then redirects the user to a Google Firebase Application hosted on the domain na2signing[.]web[.]app. The na2signing[.]web[.]app HTML subsequently redirects the user to a WordPress site, which is running yet another redirector at eaglelodgealaska[.]com.

After this final redirect, the user is sent to the attacker-controlled docusigning[.]kirklandellis[.]net website.

Trend to watch: Multi-channel phishing can start with a "benign" link

We are seeing more attacks targeting users through multiple communication channels — usually first with a link. **We refer to this attack type as multi-channel phishing**. And, according to our commissioned survey conducted by Forrester Consulting, **89% of security decision-makers are concerned about multi-channel phishing threats**[5]:

About 8 in 10

said their firm is exposed across various channels such as IM/cloud collaboration/ productivity tools, mobile/ SMS, social channels.

Only 1 in 4

respondents felt their firms are completely prepared for phishing threats across various channels.

Attack definitions

- **Multi-channel attack**
  A phishing attack that attempts to exploit a user by engaging them across multiple applications
- **Multi-vector attack**
  Attempting to gain unauthorized access by simultaneously attacking multiple entry points
- **Multi-mode attack**
  The various stages of an attack lifecycle as an attacker progresses towards their end goal

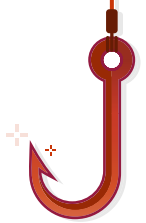

Table of contents

Trend to watch: Multi-channel phishing can start with a “benign” link

One example of multi-channel phishing involves a “deferred” attack, where the link is still benign when the email is first sent. For example:

Setup:

The attacker sets up infrastructure (such as registering a domain, setting up email authentication, and creating a benign webpage) for their future phishing attempt. At this point, email systems won’t pick up evidence of an attack.

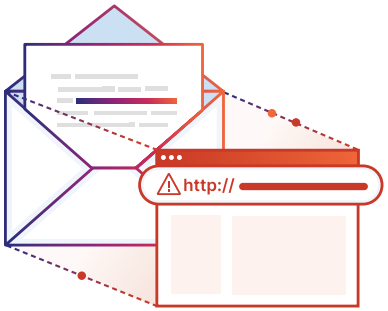

Attack launch part 2, Sunday evening:

After email delivery, the webpage is “weaponized”, for example, by updating it to include a fake login page for harvesting credentials.

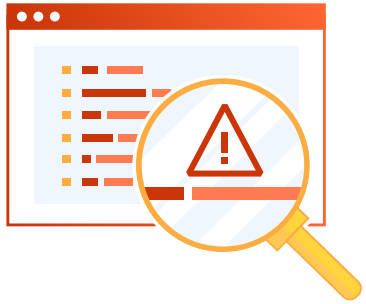

Weeks before launch | Sunday | Monday

Attack launch part 1, Sunday morning:

The attacker sends an email from the newly-created domain with a link to the still-benign webpage. Email systems don’t flag it as suspicious.

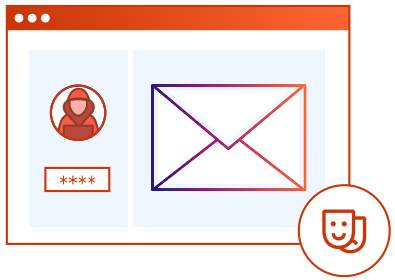

Attack landing:

Employees beginning their work week see the email. It only takes one to click and enter their credentials for the attack to succeed.

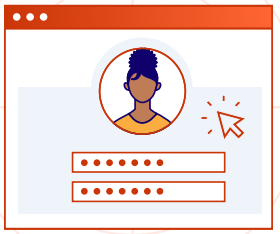

Table of contents

Look closer: A multi-channel phishing attempt on Cloudflare

In July 2022, the Cloudflare Security team **received** reports of employees receiving legitimate-looking text messages pointing to what appeared to be a Cloudflare Okta login page.

The text messages pointed to an official-looking domain (cloudflare-okta[.]com) that had been **registered less than 40 minutes** before the phishing campaign began.

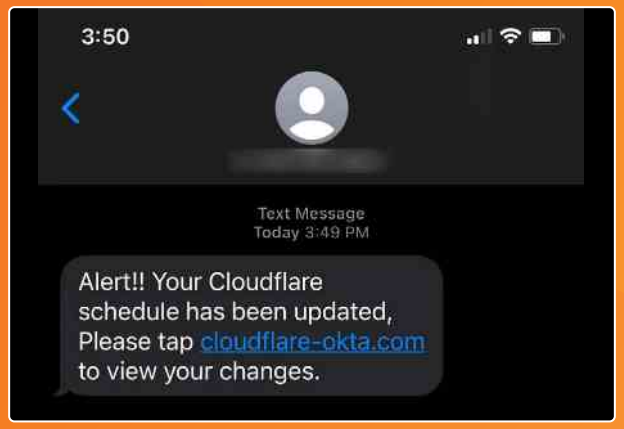

Text message received by Cloudflare employees pointing to fake Okta login page

If someone clicked on the link, it would take them to a phishing page that looked identical to a legitimate Okta login page (Cloudflare uses Okta as our identity provider), prompting visitors for their credentials.

Ultimately, if an intended victim had made it past the steps of entering credentials and a Time-Based One Time Password (TOTP) code on the phishing site, the phishing page then initiated a download of a phishing payload which included AnyDesk's remote access software. That software, if installed, would allow an attacker to control the victim's machine remotely.

However, Cloudflare does not use TOTP codes (instead, every employee is issued a physical FIDO2-compliant security key). The attackers did not get past our hard key requirement or our **SASE** platform, Cloudflare One.

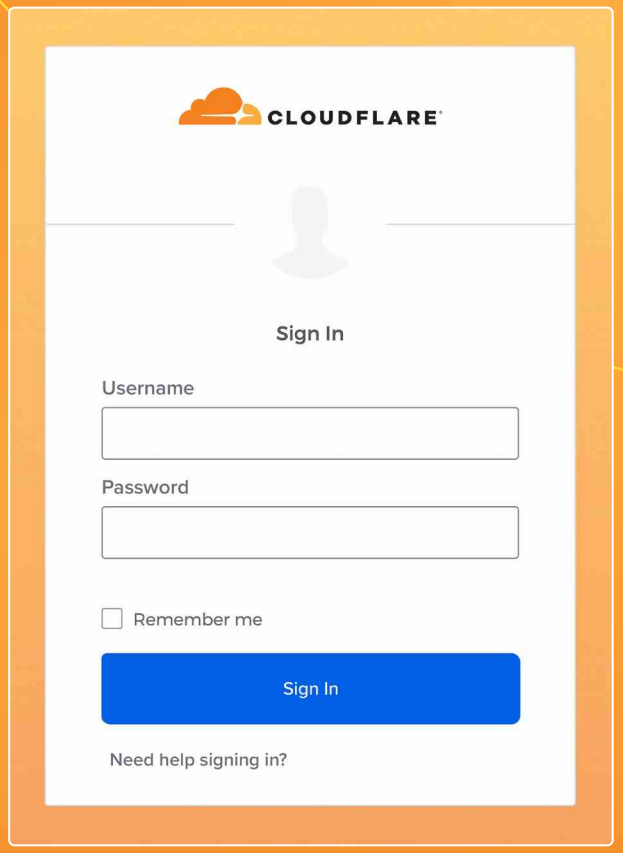

Fake Cloudflare Okta login page, nearly identical to real version

Top Threat: Identity deception

More attacks now use **identity deception (impersonating someone else's identity)** — the third-most prevalent email threat category. We observed identity deception in **14.2% of detections** from May 2, 2022 - May 2, 2023, a jump from 10.3% from the year prior[6]. This attack type takes many forms, including **brand impersonation** and **business email compromise (BEC)**.

Major challenges organizations face in preventing phishing are the volume of attacks and the difficulty of distinguishing legitimate emails and websites from fraudulent ones.

Whether they're attempting large-scale campaigns or highly targeted account compromise, **today's attackers will find ways to exploit the trust many people place in messages from "known" senders.**

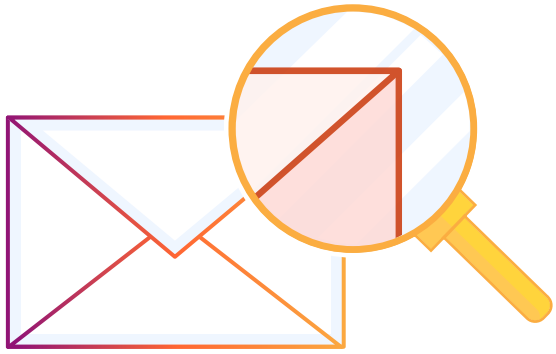

Look closer: Phishing against democracy

In the three months leading up to the 2022 US midterm elections, Cloudflare prevented around 150,000 phishing emails from reaching campaign officials. This included a phishing attempt targeting a US congressional candidate's staff members.

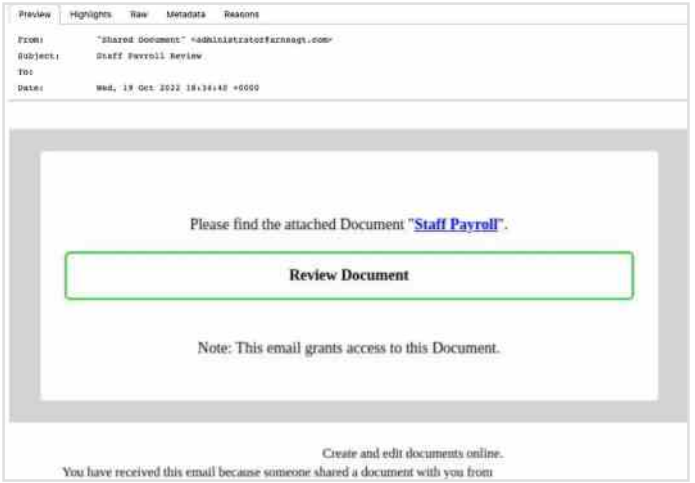

The staffers were sent an email with the subject "Staff Payroll Review," asking them to access a document link. The email contained a valid email footer and branding consistent with the campaign.

However, our models found several discrepancies within the email's metadata, including a link to a newly registered domain.

Our system blocked the email and prevented the potential loss of data and money.

Table of contents

Trend to watch: BEC threats spike worldwide

By now, many organizations have heard of business email compromise (BEC) — a specific form of financially-motivated phishing. Yet BECs continue to inflict major pain.

Why? BECs don't rely on deceptive links or malicious attachments; instead, **they exploit deep understanding of the recipient's email behaviors and business processes**. That knowledge can extend to compromising the target's trusted supply chain and partners as well.

For example, account compromise, which may be used in BEC attacks, is when the attacker actually takes control of a user's email account. If it's a partner's email account, it's also referred to as vendor email compromise (VEC). Compromised accounts can target others, since the source doesn't change.

Imagine a vendor you've trusted for a while: you email regularly about projects and even their weekend plans. Then one day, you pay a "fake" invoice — one that looks like past invoices in every way, except for a routing number change. That's because an attacker has been "inside" your email account for *weeks or even months*.

Although BEC threats represented a low volume (0.5%) of our total detections[7], we believe this is due to our technologies identifying these sooner in the attack cycle (for example, before an attacker has a chance to send a fraudulent invoice diverting payments).

Organizations that fail to thwart BEC will face more financial losses than ever before:

Over $50 billion in losses

Total domestic and international losses from BEC totaled over $50 billion from Oct. 2013 - Dec 2022[1]

17% increase

There was a 17% increase in identified global exposed BEC losses from Dec. 2021-Dec. 2022[1]

Costlier than ransomware

There were 2,385 ransomware complaints with losses of more than $34.3 million, compared to 21,832 BEC complaints with losses of over $2.7 billion, in 2022[11]

71% of organizations

71% of organizations experienced an attempted or actual BEC attack in 2022[12]

Top threat: Brand impersonation

Table of contents

With our sample dataset, attackers posed as nearly **1,000 different organizations** in nearly a billion impersonation attempts against Cloudflare customers.

However, the majority **(51.7%)** of the time, they posed as one of just 20 organizations noted below — with **Microsoft**[4] topping the list. Attackers not only frequently impersonate Microsoft, they will also use Microsoft's own tools to commit fraud.

Most impersonated brands:

1. Microsoft
2. World Health Organization
3. Google
4. SpaceX
5. Salesforce
6. Apple
7. Amazon
8. T-Mobile
9. YouTube
10. MasterCard
11. Notion.so
12. Comcast
13. Line Pay
14. MasterClass
15. Box
16. Truist Financial Corp
17. Facebook
18. Instagram
19. AT&T
20. Louis Vuitton

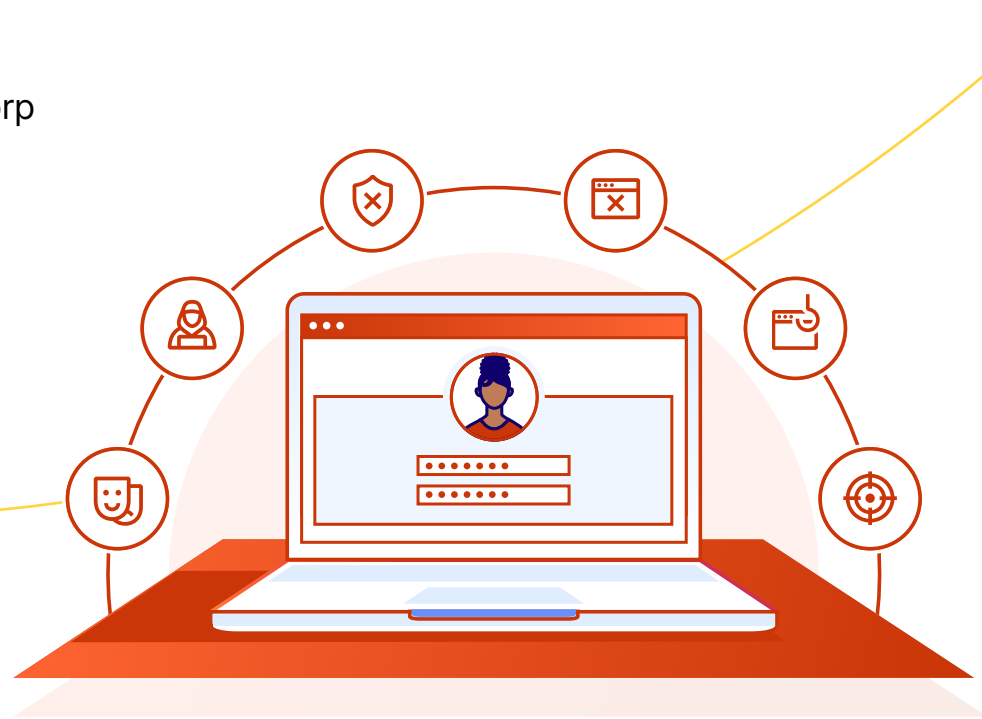

Look closer: Brand impersonation across key industries

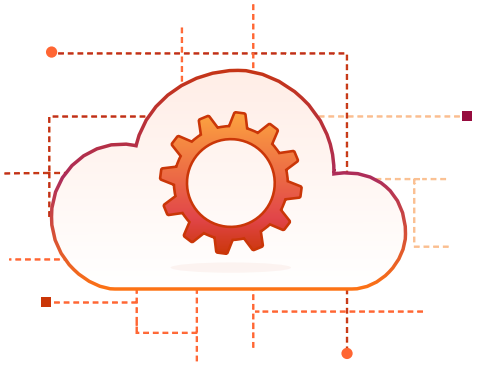

Most impersonated SaaS brands:

1. Salesforce
2. Notion.so
3. Box
4. 1Password
5. Zoom
6. Rapid7
7. Marketo
8. ServiceNow
9. NetSuite
10. Workday

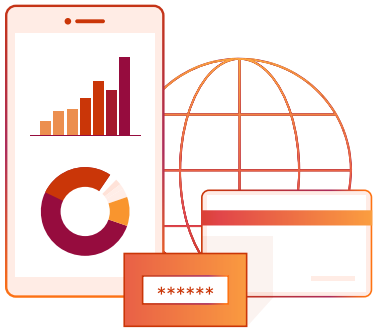

Most impersonated financial services brands:

1. MasterCard
2. Truist Financial
3. Investec
4. Generali Group
5. Bitcoin
6. OpenSea
7. Bank of America
8. Binance
9. Visa
10. Nationwide

Most impersonated social media brands:

1. YouTube
2. Facebook
3. Instagram
4. WhatsApp
5. Pinterest
6. Parler
7. Twitter
8. LinkedIn
9. Discord
10. Reddit

Based on volume of brand impersonation indicators in emails observed by Cloudflare's Area 1 email security service from May 2, 2022 to May 2, 2023

Table of contents

Look closer: Brand impersonation across the globe

Most impersonated EMEA brands:

1. World Health Organization
2. Louis Vuitton
3. Investec
4. Chanel
5. Generali Group

Most impersonated APAC brands:

1. LINE
2. JCB Global
3. State Bank of India
4. Toyota
5. Toshiba

Most impersonated LATAM brands:

1. Banco Bradesco
2. Atento
3. LATAM Airlines
4. California Supermercados
5. Locaweb

Based on volume of brand impersonation indicators in emails observed by Cloudflare's Area 1 email security service from May 2, 2022 to May 2, 2023

Trend to watch: Brand impersonation tricks common email defenses

Brand impersonation can be partially addressed with email authentication, but these pose many limitations. For example, attackers can easily configure their emails to pass authentication standards.

> In fact, **the majority (89%) of unwanted messages "passed"** SPF, DKIM, and/or DMARC checks.[8]

Here are examples of ways threats such as brand impersonation can bypass email authentication:

- **SPF, DKIM, and DMARC do not prevent** look-alike email, domain or display name spoofing
- **DKIM does not protect against** replay attacks, a type of attack that can "fool" a network into thinking the message has passed protocols
- **DMARC does not prevent spoofing** of another organization's domain
- **These standards do not inspect content** — they only determine whether the sender's domain is properly configured

Another way attackers can successfully deliver a brand impersonation or other type of phishing email is to use a **newly registered domain** (NRD). Thousands of domains are registered every day[13] — the multi-channel phishing campaign targeting Cloudflare (described earlier) was just one example of an unsuccessful attack using an NRD.

Threats due to **"Domain Age"** (in combination with a variety of other data points) **represented the #2 threat category**, detected in **30% of all unwanted emails**. *(For this report's purposes, we considered NRDs to be registered or with a change of ownership within 48 hours prior to the email send).*

Phishing attacks of all types have grown in sophistication, so much so that traditional approaches against them are not sufficient to prevent the most dangerous attacks.

To stay ahead of these and other advanced threats, read on for our recommendations.

Recommendations

1 Secure email with a Zero Trust approach

Despite email's pervasiveness, many organizations still follow a "castle-and-moat" security model that trusts messages from certain individuals and systems by default.

With a Zero Trust security model, you trust no one and nothing. No user or device has completely unfettered, trusted access to all apps — including email — or network resources. This mindset shift is especially critical if you have multi-cloud environments and a remote or hybrid workforce.

Don't trust emails just because they have email authentication set up, are from reputable domains, or "from" someone with whom you have a prior communication history. Choose a cloud email security solution rooted in the Zero Trust model and make it more difficult for attackers to exploit existing trust in "known" senders.

2 Augment cloud email with multiple anti-phishing controls

A multi-layered defense can preemptively address high-risk areas for email exposure, including:

- Blocking never-before-seen attacks in real time, without needing to "tune" a SEG or wait for policy updates
- Exposing malware-less financial fraud such as VEC and supply chain phishing
- Automatically isolating suspicious links or attachments in email
- Identifying and stopping data exfiltration, particularly via cloud-based email and collaboration tools
- Discovering compromised accounts and domains attackers use to launch campaigns

More organizations are choosing a layered approach to phishing protection. As noted in The Forrester Wave™: Enterprise Email Security, Q2 2023, "The email security vendors you work with should demonstrate an ability to connect and share data with each other and with key tools in your security tech stack."[15]

3 Adopt phishing-resistant multi-factor authentication

Any form of multi-factor authentication (MFA) is better than none, but not all MFA provides the same level of security. Hardware security keys are among the most secure authentication methods for preventing successful phishing attacks; they can protect networks even if attackers gain access to usernames and passwords. Consider replacing MFA methods like SMS or time-based OTP with more proven methods like FIDO-2 compliant MFA implementations.

Applying the principle of least privilege can also ensure hackers who make it past MFA controls can access only a limited set of apps, and partitioning the network with microsegmentation can prevent lateral movement and contain any breaches early.

4 Make it harder for humans to make mistakes

The larger your organization, the more each of your teams will want to use their own preferred tools and software. Meet employees and teams where they are by making the tools they already use more secure, and preventing them from making mistakes.

For example, email link isolation, which integrates email security with remote browser isolation (RBI) technology, can automatically block and isolate domains that host phishing links, instead of relying on users to stop themselves from clicking.

5 Establish a paranoid, blame-free culture

Encouraging an open, transparent “see something, say something approach” to collaborating with your IT and security incident response teams 24/7 helps get everyone on “team cyber.”

Minutes matter during attacks. Establishing a paranoid but blame-free culture that reports suspicious activity — as well as genuine mistakes — early and often helps ensure incidents (no matter how rare) are reported as soon as possible.

Learn more

To discover which phishing attacks your current email security systems are missing, request a complimentary phishing risk assessment. The assessment requires no hardware or software installation, and will not impact email flow.

To learn more about how Cloudflare provides Zero Trust protection against email threats, visit us here.

Appendix — Threat Glossary

Account compromise — When an attacker takes control of a user's email account. This is also referred to as *Email Account Compromise (EAC)*, which is a close relative of *Business Email Compromise (BEC)*. Attackers use a wide array of techniques such as dictionary brute forcing, credential harvesting attacks, and credential theft. The essential details are that a user's email account credentials become compromised through malicious actions. Subsequently, the attacker uses that account to send malicious content to new targets.

ASN reputation — The overall score assigned to an Autonomous System Number (ASN) based on behavior. For example, ASNs from which high volumes of spam or malicious emails originate, will tend to have poorer reputations and thus lower scores. ASNs with low reputation scores are often used in attacks.

Attachment — Any file attached to an email that, when opened or executed in the context of an attack, includes a call-to-action (e.g., lures target to click a link) or performs a series of actions set by an attacker. If the intended victim opens an attachment or clicks a malicious attachment link, they may ultimately install a piece of malware that could lead to ransomware or follow-on operations through backdoors and RATs.

Brand impersonation — A form of *identity deception* where an attacker sends a phishing message that impersonates a recognizable company or brand. Brand impersonation is conducted using a wide range of techniques. A common one is *display Name Spoofing*, where the sender display name in the visible email headers includes a legitimate brand. In addition, attackers might use *domain impersonation*. In this case, the attacker registers a domain that looks similar to the impersonated brand's domain, and uses it to send phishing messages.

Attackers often use various forms of obfuscation, such as homograph spoofing, in brand impersonation attacks. They might also register the exact same domain name as that used by the impersonated brand but with a different top level domain (TLD). These techniques can be leveraged throughout all sections of an email, including the sender display name, email address (including the sender domain name), subject line, body content (HTML and plaintext), hypertext for links, and hyperlinks themselves (i.e., the actual URLs).

Business email compromise (BEC) — An increasingly common, effective, and costly targeted email attack designed to trick recipients into transferring funds, typically through forged invoices, to scammer accounts. BEC falls into various categories based on its sophistication, ranging from using a spoofed email to compromising a vendor in a supply chain attack.

Credential harvesters — Sites set up by an attacker to deceive users into providing their login credentials. This particular attack presents the user with a page that imitates an email or other account login page. Unwitting users may enter their credentials, ultimately providing attackers with access to their accounts. Because people often reuse passwords for multiple accounts, a member of your organization providing credentials to a harvester may give an attacker access to many accounts.

Domain reputation (related to *Domain Age*) — The overall score assigned to a domain. For example, domains that send out a large number of new emails immediately after domain registration will tend to have a poorer reputation, and thus a lower score. Whereas older, known domains tend to have a positive reputation, and thus a higher score. Domains with low reputation scores are often used in attacks.

Extortion — This tactic is commonly used to force a person or organization to perform a set of actions they would not otherwise normally perform. This is typically done under duress; for example, asking the intended victim to pay a ransom during a DDoS attack. The level of extortion can lead to a wide range of compromise depending on the attacker's intentions and resources.

Identity deception — This occurs when an attacker or someone with malicious intent sends an email claiming to be someone else. The mechanisms and tactics of this vary widely. Some tactics include registering domains that look similar (aka *domain impersonation*), are *spoofed*, or utilize display name tricks to appear to be sourced from a trusted domain. Other variations include sending email using domain fronting and high-reputation web services platforms.

Link — When clicked, a deceptive link will open the user's default web browser and render the data referenced in the link, or open an application directly (e.g. a PDF). Since the display text for a link (i.e., hypertext) in HTML can be arbitrarily set, attackers can make a URL appear as if it links to a benign site when, in fact, it is actually malicious. Malicious links can lead to arbitrary code execution or Remote Code Execution (RCE), credential harvesting, click fraud, unwanted installs, and other compromises.

Scam — A broad category of phishing fraud. The foundation is to entice a victim to provide money under a promise of a product, service, good, or even significant sum of money in return. The common theme is the transfer of money in a method that is atypical for the sender. Changes in common payment practices or sudden demands to pay sums via wire transfer can also be indicators.

Voice phishing — Also called "vishing," this usually refers to the practice of leaving fake voice messages in hopes that victims will call back to provide personal information (such as bank and credit card details), which will be used in other attacks. In our email security detections, we have observed attackers combining email and voice vectors by sending emails with attachments of a voicemail recording, media file or a link to a file. We have also observed attackers sending emails that had no malicious payloads, just a phone number.

Other — For the purpose of this report, other threat indicator categories with statistically insignificant numbers have been consolidated into the "other" category. This includes command and control (any attempt to launch a process on a host system), IP policy (detection based on a customer-specific policy), target development (attacker information-gathering to facilitate a successful attack), among others.

Endnotes

[1] “Business Email Compromise: The $50 Billion Scam.” IC3.gov, June 9, 2023. https://www.ic3.gov/Media/Y2023/PSA230609

[2] Based on a sample of threat indicators (“categories”) detected by the Cloudflare email security service between May 2, 2022 - May 2, 2023. These indicators lead to email dispositions of malicious, BEC, spoof, or spam. Individual messages may contain multiple threat categories such as “Identity Deception”, “Brand Impersonation”, “Link”, and others that are described in the appendix.

[3] Based on messages categorized as either “Malicious” or “Malicious-BEC” by the Cloudflare email security service between May 1, 2022 - April 30, 2023.

[4] Based on an aggregate volume of brand impersonations (see appendix) observed by the Cloudflare email security service between May 2, 2022 - May 2, 2023. For our analysis: “Microsoft” brand impersonations also included impersonations of “Windows”, “Outlook”, “Office365”, “Microsoft Teams” “Windows Defender”, “SharePoint”, “Yammer”, “OneDrive”, Skype”, and “OneNote”; “Google” brand impersonations also included impersonations of “Gmail”, and “Hangouts”; “Amazon” brand impersonations also included impersonations of “Amazon Fresh”; “Apple” brand impersonations also included impersonations of “iTunes” and “iCloud”; and “Salesforce” brand impersonations included impersonations of “ExactTarget.”

[5] Source: Forrester Opportunity Snapshot: A Custom Study Commissioned by Cloudflare, “Leverage Zero Trust to Combat Multichannel Phishing Threats,” May 2023. *Methodology: This Opportunity Snapshot was commissioned by Cloudflare. To create this profile, Forrester Consulting supplemented existing Forrester research with custom survey questions asked of 316 global practitioners at the manager level or above who are responsible for their organizations’ security strategy. The custom survey began in January 2023 and was completed in February 2023.*

[6] Based on a sample of threat indicators categorized as “IdentityDeception” (see appendix) by the Cloudflare email security service between May 2, 2022 - May 2, 2023, and as “IdentityDeception” between May 1, 2021 - April 30, 2022 by Area 1 Security. Area 1 Security was acquired by Cloudflare in April 2022.

[7] Based on a sample of threat indicators categorized as “BEC” or “BECType1” (see appendix) by the Cloudflare email security service between May 2, 2022 - May 2, 2023, and “BEC” or “BECType1” between May 1, 2021 - April 30, 2022 by Area 1 Security. Area 1 Security was acquired by Cloudflare in April 2022.

[8] Based on a sample of messages given a disposition of “malicious,” “BEC,” “spoof”, or “spam” by the Cloudflare email security service between May 2, 2022 - May 2, 2023, that also passed SPF, DKIM, and/or DMARC email authentication checks.

[9] Based on a sample of threat indicators categorized as “Links” (see appendix) by the Cloudflare email security service between May 2, 2022 - May 2, 2023, and categorized as “Links” by Area 1 Security between May 1, 2021 - April 30, 2022. Area 1 Security was acquired by Cloudflare in April 2022.

[10] “2023 Verizon Data Breach Investigations Report (DBIR).” Verizon.com, last accessed 15 June 2023. https://www.verizon.com/business/resources/reports/dbir/

[11] “Federal Bureau of Investigation Internet Crime Report 2022.” IC3.gov, accessed 15 June 2023. https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf

[12] “2023 AFP Payments Fraud and Control Survey.” AFPonline.org, accessed 15 June 2023. https://www.afponline.org/publications-data-tools/reports/survey-research-economic-data/Details/payments-fraud

[13] “The Domain Name Industry Brief: Q1 2023 Data and Analysis.” Verisign.com, accessed 15 June 2023. https://www.verisign.com/en_US/domain-names/dnib/index.xhtml?section=executive-summary

[14] Based on a sample of threat indicators categorized as “DomainAge” (see appendix) by the Cloudflare email security service between May 2, 2022 - May 2, 2023.

[15] Source: Forrester Research, “The Forrester Wave™: Enterprise Email Security, Q2 2023,” June 12, 2023. *The Forrester Wave™ is copyrighted by Forrester Research, Inc. Forrester and Forrester Wave are trademarks of Forrester Research, Inc. Forrester does not endorse any vendor, product, or service depicted in the Forrester Wave. Information is based on best available resources. Opinions reflect judgment at the time and are subject to change.*

Table of contents

Call: 1 888 99 FLARE
Email: enterprise@cloudflare.com
Visit: www.cloudflare.com

REV:BDES-4838.2023AUG14