egress

northdoor
◦ Store IT ◦ Protect IT ◦ Use IT ◦

October 2023

Phishing Threat Trends Report

From pretexting to payloads, how have phishing attacks evolved in 2023?

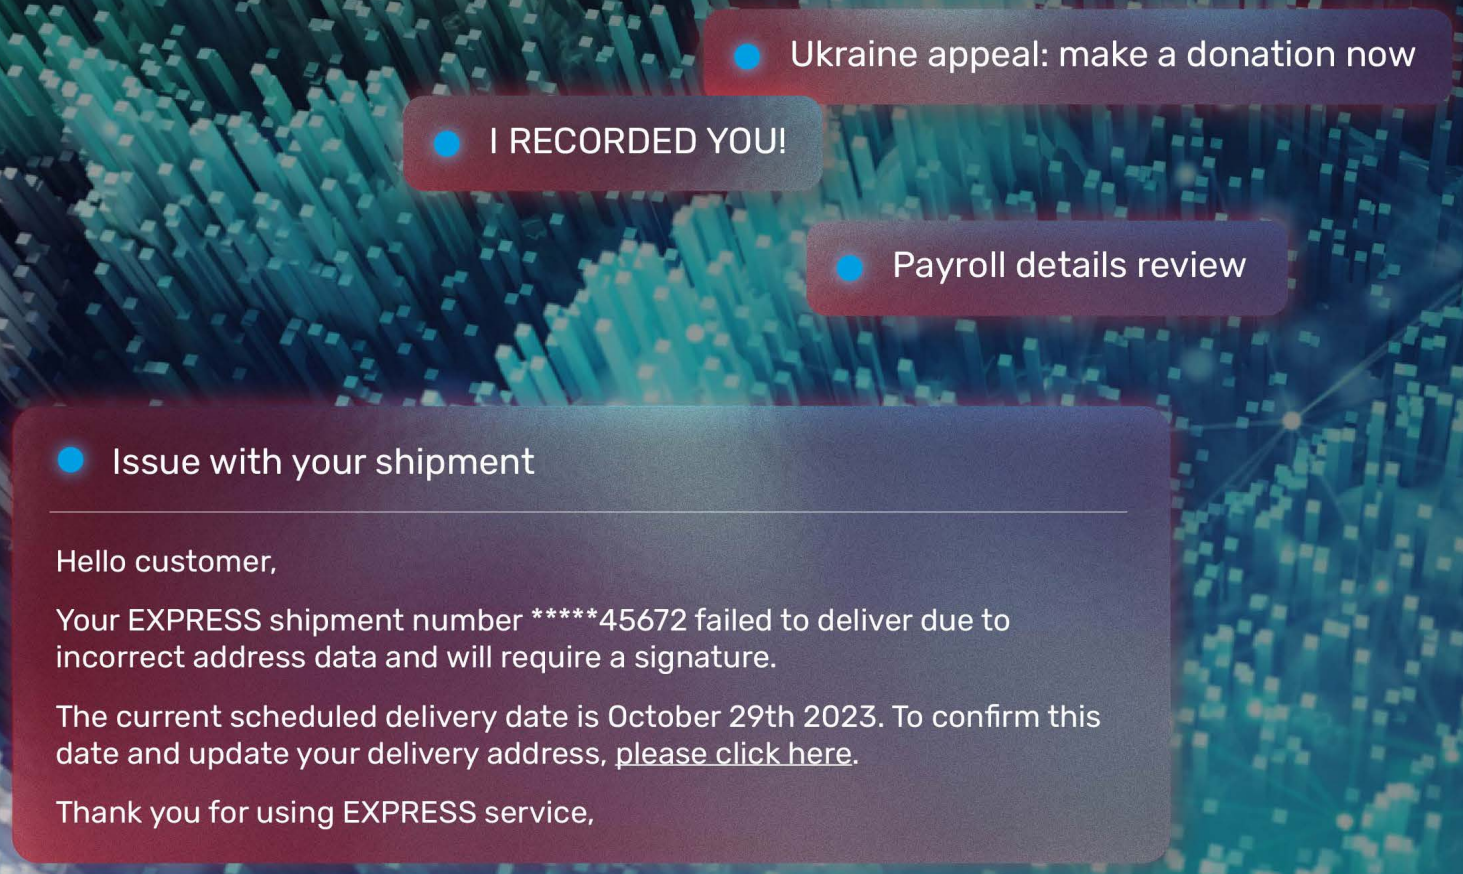

WELCOME

As we enter the final straight of 2023, our second Phishing Threat Trends Report takes a broad look at the phishing trends we've seen so far. We review the most phished topics of this year and predict what's to come next, as well as examine the prevalent obfuscation techniques that cybercriminals are using to bypass perimeter defenses. We also couldn't let this edition pass without looking at the impact of chatbots on the phishing threat landscape.

All statistics used in this report have been generated using data from Egress Defend, an integrated cloud email security solution. As always, please reach out if you have questions about this report or how Egress can improve your email security.

Jack Chapman

VP of Threat Intelligence, Egress

IN THIS EDITION

MOST PHISHED TOPICS OF 2023

JANUARY

RingCentral impersonation

Brand impersonations of RingCentral were the most common phishing emails detected in January, frequently using an HTML smuggling attack payload disguised as a missed message.

FEBRUARY

Alias impersonation attackers

One-third (32.2%) of attacks detected were alias impersonation, frequently leveraging social engineering. This could be related to the use of large language models to generate attacks.

MARCH

HMRC/IRS notification impersonation

With taxes needed to be filed by April 18, 2023, and fiscal years starting in many countries globally, attackers mimicked organizations such as the IRS and HMRC and used upcoming deadlines to pressure their targets.

APRIL

Security software impersonation

15.2% of attacks impersonated security software companies, leveraging trusted brands such as Avast and Norton Security.

MAY

Sextortion phish/life ruiners

'Life ruining' phishing emails, such as sextortion attacks, were the most common in May. Cybercriminals claim to have leverage over the recipient, who is manipulated into an emotional response.

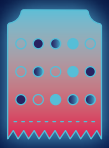

JUNE

Lottery impersonation

12.7% of attacks used fraudulent lottery and other related topics to get targets to click phishing hyperlinks and share their bank details with attackers.

JULY

Salesforce/Meta ads

Brand impersonation continued in July, with attacks pretending to come from Meta and Salesforce. This is related to a zero-day vulnerability in Salesforce's email services and SMTP servers, which was used to launch a phishing campaign targeting Facebook accounts.

AUGUST

Geek Squad

A popular topic throughout 2023, and first detected in May, Geek Squad phishing emails surged again in August.

SEPTEMBER

Credit card payments

In September, 15.1% of phishing attacks came from cyber criminals who impersonated well-known banks or failed credit card transactions. They then used this sense of urgency to scrape their victims credentials.

PREDICTIONS FOR REST OF THE YEAR

OCTOBER

Fax impersonation attacks

Fax-related attackers rose sharply in October 2022, at 30.8% of detected phishing attacks. This is a tried and true attack that cybercriminals can use throughout the year – and potentially we'll see them rise again this year.

NOVEMBER

Crypto scams / Black Friday attacks

Crypto phishing attacks that attempt to trick victims into giving up their cryptocurrency keys may increase in November, alongside Thanksgiving and Black Friday themed attacks.

DECEMBER

Christmas related

Christmas remains a top phishing topic in December, with many brand impersonation attacks offering fraudulent discounts and competition prizes. The most impersonated brands in December 2022 included Amazon, Walmart, and PayPal – and it's likely these three household names will be leveraged again this year.

Most phished topic of the year - 2023

Missed voice message

Missed voice messages accounted for 18.4% of phishing attacks between January to September 2023, making them the most phished topic for the year so far. Many of these attacks use HTML smuggling to hide their payload – a topic that we unpack further on pages 10-12.

HAVE CHATBOTS REALLY REVOLUTIONIZED PHISHING?

The potential for cybercriminals to use large language models has dominated headlines in 2023 – but has it really changed the game?

The launch of ChatGPT in November 2022 threw fuel on the fire of the AI race, and since then, the news has been saturated with the latest developments and applications, not least in the cybersecurity industry.

Like many innovations, the power of large language models (LLMs) lies in the hands of those who wield them – and the potential for cybercriminals to use chatbots to create phishing campaigns and malware has been cause for concern. In our 2023 Email Risk Report, 72% of cybersecurity leaders stated they were worried about the use of chatbots to improve phishing attacks.

But how likely is this in reality?

LLMs lower the barrier for entry to cybercrime

Once upon a time, it was possible to spot a high proportion of phishing emails from their poor grammar, bad spelling, and ridiculous requests. Those attacks were far less sophisticated before the launch of ChatGPT, and have increased in number due to the commoditization of cybercrime, the growth of related criminal networks, and the crime-as-a-service ecosystem that includes offerings such as phishing kits and pre-written malware.

LLMs lower the barrier for entry further, making it possible to create well-written phishing campaigns and generate malware that less capable coders couldn't produce alone, and which can be further refined depending on their coding ability. This has the greatest impact at the 'lower end' of the spectrum, such as for less experienced cybercriminals, those with limited coding ability, or those who don't write fluently in the target language.

Probably the most concerning but least talked about application of LMMs is reconnaissance for highly targeted attacks. Within seconds, a chatbot can scrape the internet for open-source intelligence (OSINT) about a chosen target that can be leveraged for social engineering.

Additionally, LLMs can increase the velocity of attack, creating a higher volume of attacks far more rapidly than a person can (Figure 1).

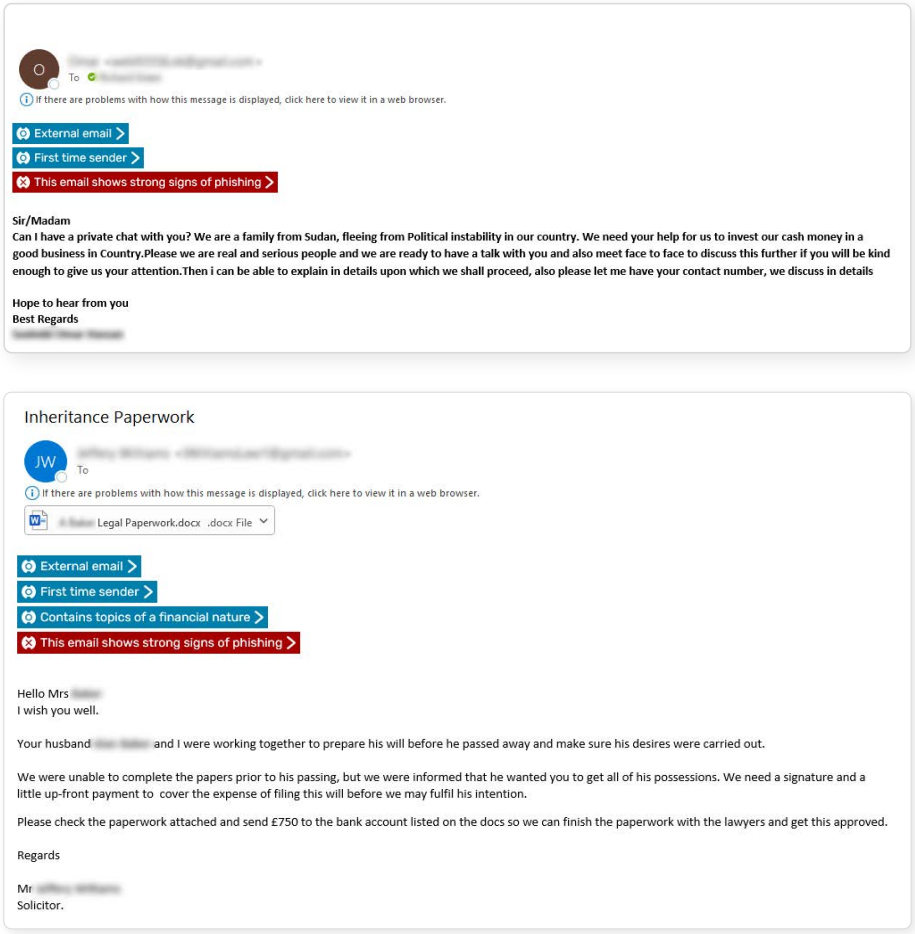

FIG 1. THESE TWO EXAMPLES SHOW THE EVOLUTION OF A 419 SCAM. THE FIRST IS A RATHER OBVIOUS AND WELL-KNOWN ATTACK THAT MOST RECIPIENTS WOULD IDENTIFY. THE SECOND INVOLVES DETAILED PRETEXTING TARGETING A VULNERABLE RECIPIENT AND ASKING FOR A SMALL UPFRONT PAYMENT. IT IS NOT POSSIBLE TO CONCLUSIVELY SAY WHETHER THE SECOND EXAMPLE WAS WRITTEN BY A CHATBOT, BUT THE STYLE AND LANGUAGE REFLECT TESTS CARRIED OUT BY OUR THREAT INTELLIGENCE TEAM.

Pretexting vs. payloads: How common are social engineering emails?

Our Threat Intelligence analysts have correlated the character lengths of phishing emails with its payload. For this, they analyzed 1.7 million phishing emails to determine whether they contained either a phishing link or attachment-based payload, or whether they were 'payloadless' and relied solely on social engineering.

Phishing emails with fewer than 100 characters are 90% more likely to contain an attachment as a payload. Links to phishing websites are the most prevalent between 100 – 1,199 characters. At 1,500 characters or more (between approximately 200 - 375 words), the attack is more likely to rely on social engineering.

There's also a 'baton pass' between attachments and social engineering between approximately 500 – 1,300 characters, where fraudulent attachments and social engineering combine as part of sophisticated business email compromise (BEC) attacks.

It makes sense – the longer the attack, the more likely it is to have complex pretext to convince the target to take the requested action (Figure 2).

FIG.2 HOW LONG IS A PHISHING EMAIL? CORRELATION BETWEEN EMAIL TEXT LENGTH AND TYPE OF PAYLOAD.

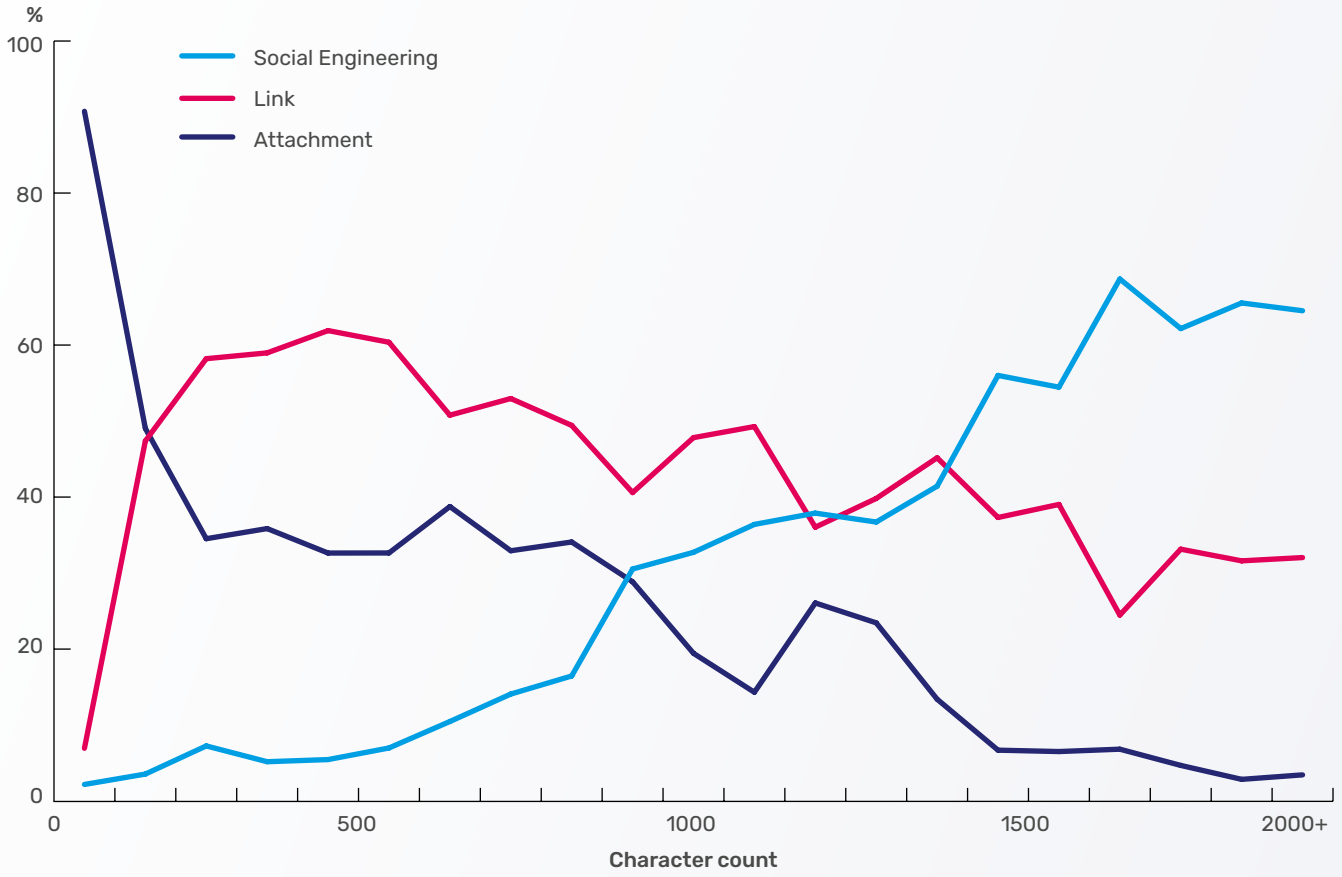

Instances of social engineering have increased slightly in prevalence so far in 2023. Analyzing phishing email received by organizations between January and December 2022, our threat analysts found that 17.98% relied solely on social engineering. Between January and September 2023, this figure rose to 19.01%. Typically, social engineering tactics are combined with a different payload (Figure 3).

FIG.3 VOLUME OF PHISHING EMAILS RELYING SOLELY ON SOCIAL ENGINEERING

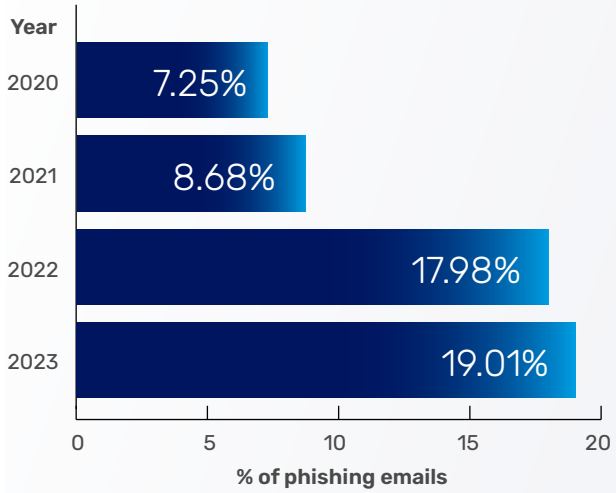

How worried should you really be about LLMs and phishing?

It all comes down to the defense you have in place. If you're relying solely on traditional perimeter detection that uses signature-based and reputation-based detection, then you urgently need to evaluate integrated cloud email security (ICES) solutions that don't rely on definitions libraries and domain checks to determine whether an email is legitimate or not.

Instead, ICES put AI in the defenders' hands, using models such as natural language processing (NLP) and natural language understanding (NLU) to analyze email content for the linguistic markers of phishing. These solutions also use machine learning to detect phishing emails with zero-day and emerging malware payloads.

Ultimately, phishing remains the main stressor – not LLMs. It doesn't matter whether an attack was written by a human or a bot – it only matters whether your defenses can detect it.

With 44.9% of phishing emails not meeting the 250-character limit and a further 26.5% falling below 500, currently AI detectors either won't work reliably or at all on 71.4% of attacks.

Ask the expert

JACK CHAPMAN,
VP OF THREAT INTELLIGENCE

Q **Is it possible to tell whether a phishing email has been written by a chatbot?**

A No. Currently, no person or tool can definitively tell you whether an attack was written by a chatbot. Because they utilize LLMs, the accuracy of most detector tools increases with longer sample sizes, often requiring a minimum of 250 characters to work. Attackers often employ obfuscation techniques such as paraphrasing chatbot output, therefore ensuring it is not identical to the original.

With 44.9% of phishing emails not meeting the 250-character limit and a further 26.5% falling below 500, currently AI detectors either won't work reliably or at all on 71.4% of attacks.

NUDGE, NUDGE...

Real-time teachable moments improve employees' ability to accurately identify phishing emails.

Legacy approaches to email security rely heavily on quarantine, barring end users from seeing or interacting with suspected phishing emails. While this approach keeps some threats away from end users, it's inevitable that they will be exposed to phishing emails at some point. When they are, you want them to be as prepared as possible to correctly identify an attack.

This is one of the reasons why Egress flips the quarantine model on its head, adding dynamic banners to neutralized threats within the inbox. These banners are designed to clearly explain the risk in a way that's easy to understand, timely, and relevant, acting as teachable moments that nudge people into improved security behaviors. This is better for users but also for administrators, saving them time and stress.

And they work.

We analyzed the emails that end users reported via Egress Defend in three of our customers between January and August 2023. As the following chart shows, people's ability to accurately identify phishing emails increased as the real-time teachable moments take effect (with a corresponding decrease in incorrectly identifying benign spam and graymail as phishing, Figure 4).

This approach augments existing security awareness training, which can be used to convey over-arching one-size-fits-all messages and meet organization's compliance requirements. At the same time, it helps to reinforce the learnings people forget or weren't listening to in the first place.

Ultimately, teaching someone to catch a phish is a more sustainable approach for long-term resilience.

FIG.4 IMPROVEMENT IN CORRECTLY IDENTIFYING PHISHING EMAILS REPORTED VIA EGRESS DEFEND

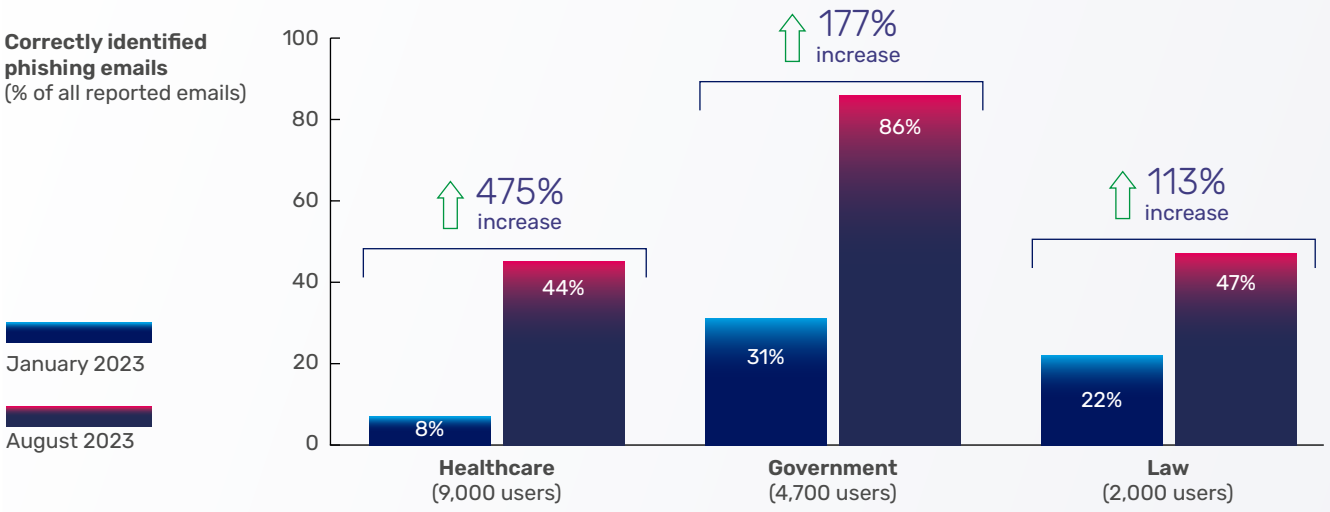

HIDING IN PLAIN SIGHT

Over half (55.2%) of phishing emails contain obfuscation techniques to help cybercriminals avoid detection.

The proportion of phishing emails employing obfuscation techniques has jumped by 24.4% so far in 2023, sitting at 55.2%. Obfuscation enables cybercriminals to hide their attacks from certain detection mechanisms.

Obfuscation techniques 101

Here's a quick overview of the techniques mentioned in this report:

Left-to-right override
Spoofing technique used to disguise attachment types or trick NLP detection within body copy.

Whitespace
Using white font on a white background to disguise characters within an email.

Homoglyphs (lookalike characters)
A form of spoofing that uses similar or identical characters or exploits UNIcode to mimic Latin characters.

Image-based
The body of the email is an image (no text is written into the email).

Hijacking legitimate hyperlinks
Attacker hosts a malicious payload on a legitimate site or uses a legitimate website link to mask the ultimate destination.

HTML smuggling
'Smuggling' an encoded malicious script within an HTML attachment.

Encoding
Content within an attachment is rendered unreadable by detection technologies.

Almost half (47.0%) of phishing emails that use obfuscation contain two layers to increase the chances of bypassing email security defenses to ensure successful delivery to the target recipient. Less than one-third (31.0%) use only one technique. Three or more layers account for the remaining 21.2% (Figure 5).

FIG.5 THE MANY LAYERS OF OBFUSCATION
PROPORTION OF PHISHING EMAILS CONTAINING DIFFERENT LAYERS OF OBFUSCATION

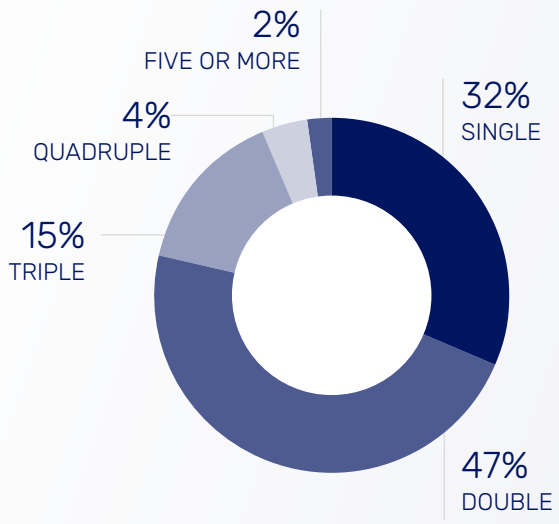

In the never-ending game of cat and mouse between cybercriminals and defenders, obfuscation techniques rise and fall in popularity as detection capabilities adapt to newer attacks. Looking at the primary obfuscation techniques in attacks sent to date in 2023, HTML smuggling has proven the most popular, accounting for 34.0% of instances. Hijacking legitimate hyperlinks is the second most popular (at 17.0%) and image-based attacks rank third (at 13.0%).

Both HTML smuggling and image-based attacks have jumped significantly in popularity since 2022: instances of image-based attacks have nearly quadrupled and HTML smuggling has almost tripled. These techniques are highly evasive as they effectively bypass traditional signature-based perimeter detection, because they hide the elements the detection capability is looking for. In HTML smuggling attacks, the payload is masked by HTML5 and JavaScript that appears benign, with the payload assembled post-delivery. Image-based attacks remove all characters from an email, leaving only a hyperlink for signature-based detection to scan. When paired with a hijacked legitimate hyperlink or a phishing website not yet listed on blocklists, the attack will get through perimeter detection.

Conversely, left-to-right override (LTRO) within the body copy of phishing emails and the use of white text on white email backgrounds have declined between January to September 2023 versus 2022. These techniques can be described as 'one and done': it is not possible to evolve them further, so once detection capabilities can identify them (which they now can) they will offer diminishing returns and decline in popularity (Figure 6).

FIG.6 CHANGES IN POPULARITY OF PRIMARY OBFUSCATION TECHNIQUES
2022 VS. 2023

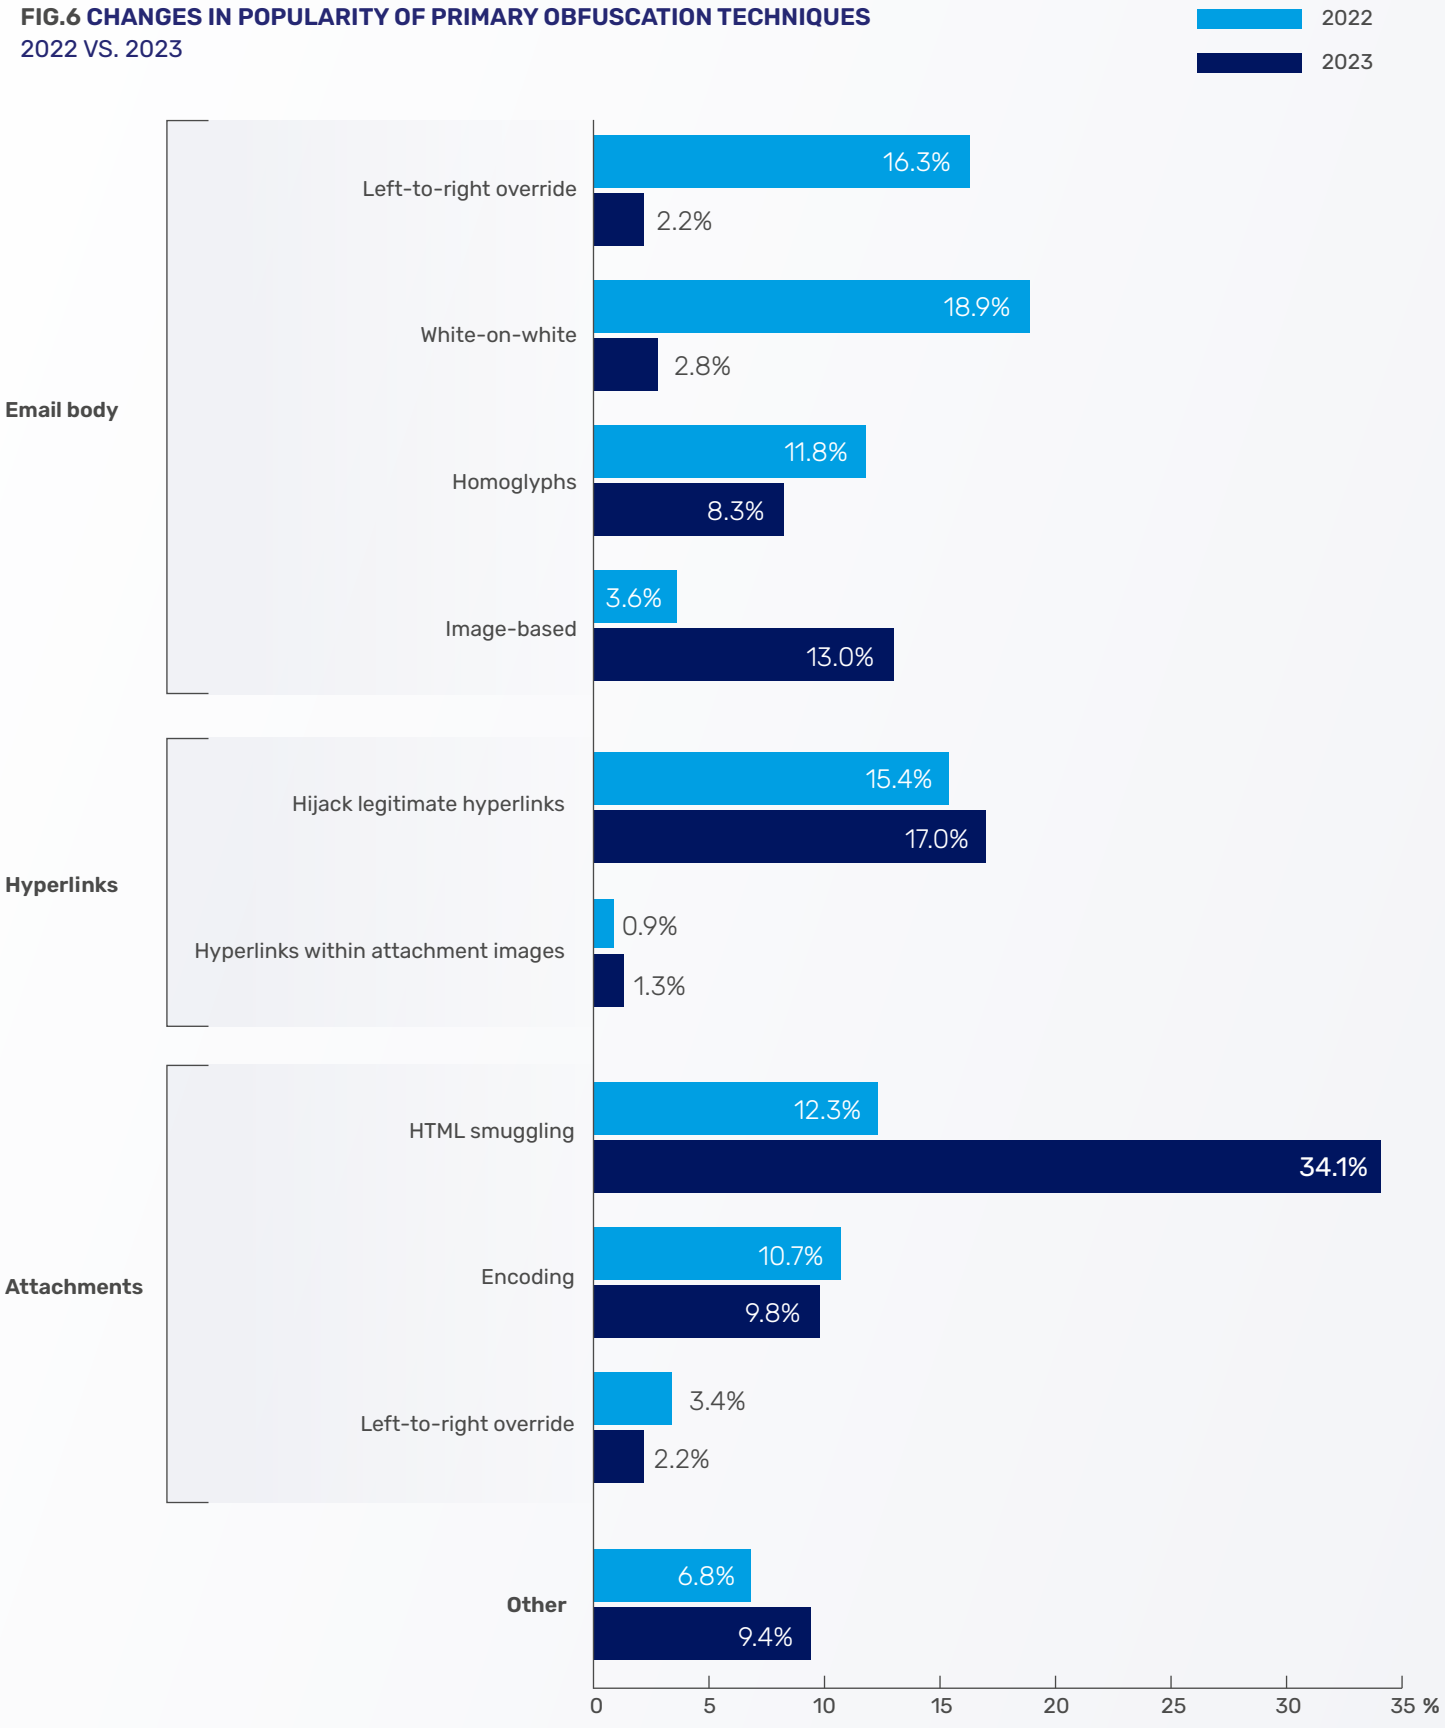

Obfuscation most used in impersonation attacks

Impersonation-based attacks are the most likely to use multiple types of obfuscation, such as 23.5% of phishing emails that impersonate a known company (such as a customer or supplier). Close behind are missed delivery attacks impersonating mail carriers (22.6%), missed voicemail messages (21.2%), and brand impersonation (19.5%) of popular organizations that are not within the supply chain (Figures 7-9).

FIG.7 DHL IMPERSONATION WHITESPACE AND NLP BREAKING ATTACK A MISSED DELIVERY PHISHING EMAIL IMPERSONATING DHL THAT WAS DETECTED BY EGRESS DEFEND. THE ATTACK CONTAINS WHITE CHARACTERS PLACED ON A WHITE BACKGROUND, LEAVING ONLY THE BLACK TEXT VISIBLE UNLESS THEY HIGHLIGHT THE BODY COPY (AS SHOWN IN THE SECOND IMAGE). THIS IS DESIGNED TO BREAK SOME NATURAL LANGUAGE PROCESSING (NLP) SYSTEMS AND BYPASS DETECTION

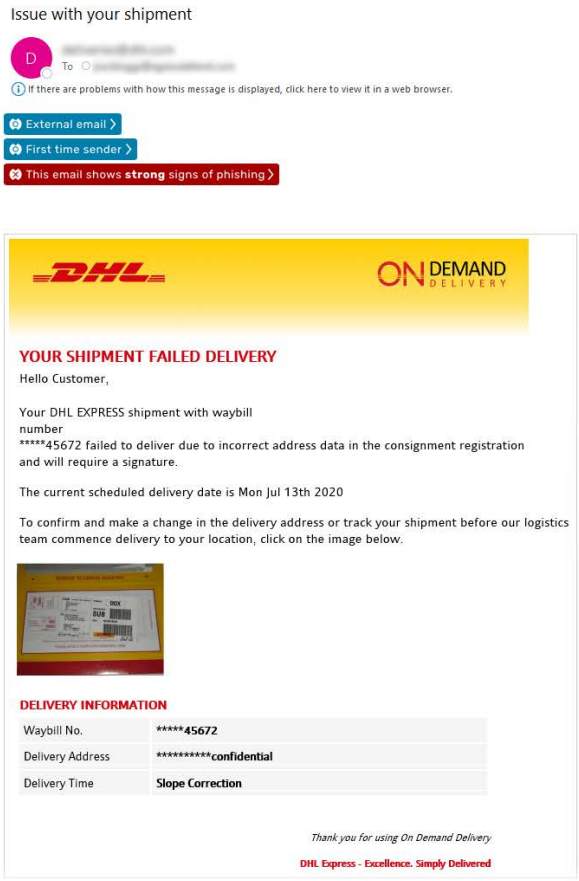

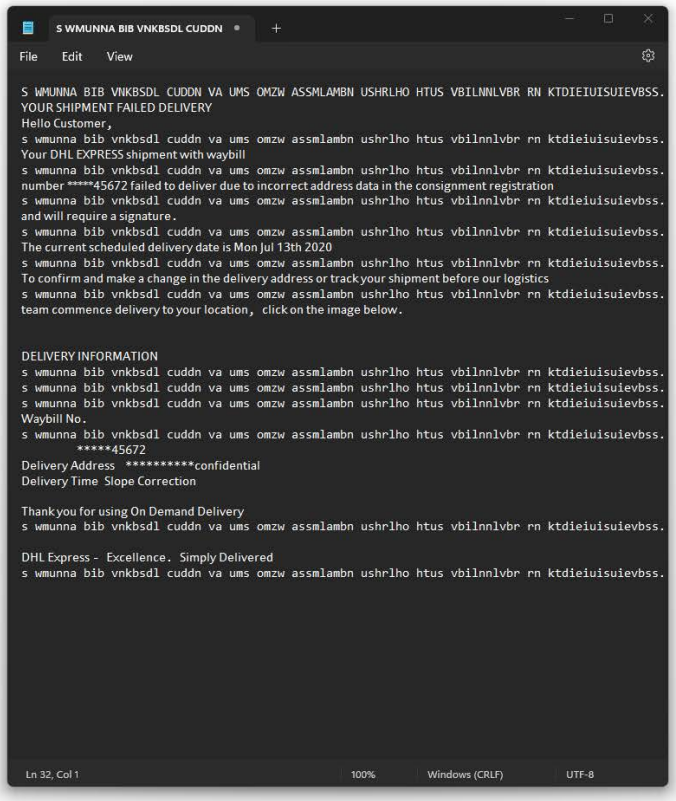

FIG.8 DOCUSIGN BRAND IMPERSONATION ATTACK THAT HIJACKED A LEGITIMATE HYPERLINK WITH THE INTENTION OF REDIRECTING THE TARGET TO A PHISHING WEBSITE

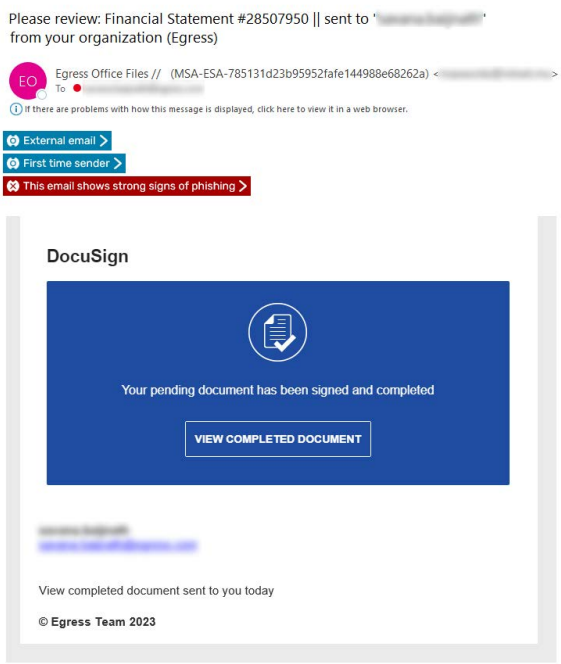

FIG.9 IMAGE-BASED ATTACK IMPERSONATING MICROSOFT. THE IMAGE WAS HYPERLINKED TO A CREDENTIAL HARVESTING SITE THAT ALSO IMPERSONATED MICROSOFT

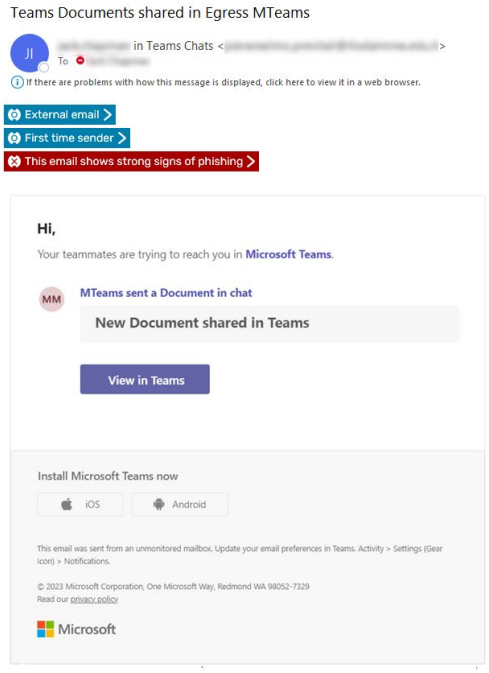

Detecting the supposedly undetectable

Obfuscation aims to make phishing emails undetectable by both traditional perimeter solutions and integrated cloud email security. Detecting them requires holistic analysis of all inbound email using intelligent technologies, without over-reliance on one mechanism – such as identifying a known payload (signature-based detection) or scanning body copy (NLP). This approach also helps to future-proof defenses against the inevitable evolution of techniques.

CATCH OF THE GRAY

At best, graymail is a distraction.
At worst, it increases your phishing risk.

To understand how graymail impacts cybersecurity, our researchers analyzed 163.8 million emails received by a subject set of organisations over a four-week period. They found that, on average, one-third (34%) of mail flow can be categorized as graymail (bulk but solicited emails such as notifications, updates, and promotional messages).

The level of graymail someone receives varies based on the industry they work in and the job they do. The highest levels of graymail are received by organizations operating in HR and recruitment, marketing, legal, finance, and education, and by individuals working in finance, HR, customer service, and office management roles.

Wednesday and Friday are the most popular days of the week to send/receive graymail. It's likely senders target these days following the traditional Monday to Friday working week, attempting to avoid the busyness of Monday and Tuesday, while trying to engage more easily distracted employees midweek and when winding down for the weekend on Friday (Figure 10).

What does graymail mean for cybersecurity?

1. More graymail = more phishing emails

Our researchers found a direct correlation between the volumes of graymail and phishing emails received. Four out of the top five industries for graymail also received the highest level of phishing emails, with education and retail as the only exceptions.

TOP FIVE INDUSTRIES BY VOLUME OF GRAYMAIL AND PHISHING EMAILS RECEIVED.

Highest volume of graymail	Highest volume of phishing
1. HR	1. HR
2. Marketing	2. Legal
3. Legal	3. Finance
4. Finance	4. Marketing
5. Education	5. Retail

FIG.10 AVERAGE GRAYMAIL VOLUME PER DAY
AVERAGE VOLUME OF GRAYMAIL AS A PERCENTAGE OF TOTAL MAIL FLOW

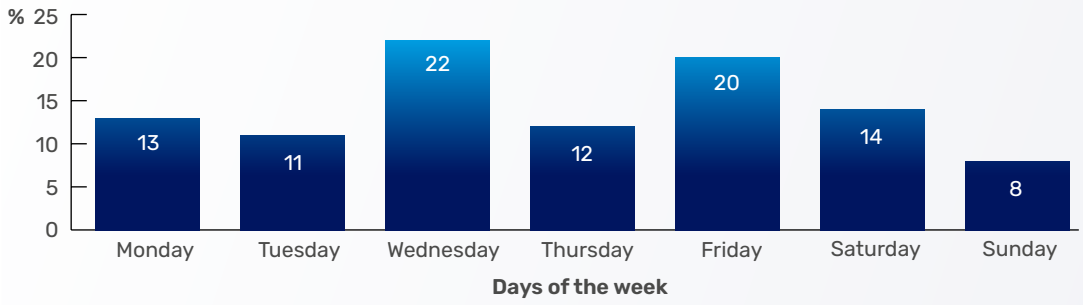

It's inevitable that organizations that rely heavily on email to communicate will have increased graymail volumes. This heavy reliance also makes them prime targets for cybercriminals looking to disguise phishing emails within busy mailboxes. In recent years there has been an increase in impersonation attacks mimicking graymail emails, such as SharePoint and social media notifications.

Additionally, any breach of client details at an organization that sends graymail could result in lists of email addresses being sold or dumped online and used in attacks.

2. Quick to click reporting floods abuse mailboxes

One of the problems with graymail is that, although solicited, it is considered a nuisance by many recipients, who may forget what they have subscribed to or be hit with more messages than anticipated.

Consequently, employees can end up flooding abuse mailboxes with graymail and other 'harmless' spam, burdening administrators with more emails to sort through and delaying them from taking action on real phishing emails that have been legitimately reported.

Quick tips to address graymail cyber risk

Filtering graymail will decrease noise within employee's mailboxes, which is not only good for efficiency but for cybersecurity too. With fewer graymail emails hitting the inbox, fewer will be incorrectly reported to the abuse mailbox meaning administrators can spend more time assessing and triaging actual phishing emails.

In addition, organizations should examine how they educate people about phishing risk. Moving beyond one-off training, dynamic banners can be added to inbound emails, using neutralized phishing emails to provide continuous, timely, and relevant education (you can see the banners used by Egress Defend in screengrabs elsewhere in this report). This approach is proven to increase people's ability to accurately report phishing emails, even when no banners are present as part of phishing simulations.

BY THE NUMBERS: A QUICK HIT OF PHISHING TRENDS IN 2023

Your questions answered with a round-up of phishing statistics.

Q Has phishing become more prevalent in 2023?

A No, but more phishing emails are getting through traditional perimeter detection, so it probably feels like overall volume has increased. As this report has shown, attacks are increasing in sophistication and cybercriminals use a multitude of tactics to get through perimeter email security.

CHANGES IN HOW PHISHING ATTACKS HAVE EVADED PERIMETER SECURITY JANUARY – SEPTEMBER 2023 VS. 2022

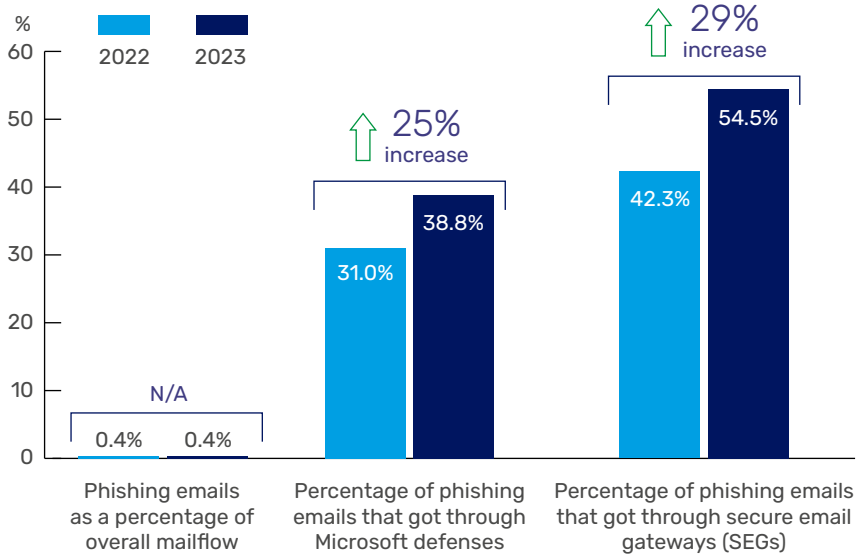

Q Are more phishing emails being sent from compromised accounts?

A Yes. There's been an 11% increase in phishing attacks sent from compromised accounts so far in 2023. Compromised accounts are trusted domains, so these attacks usually get through traditional perimeter detection. Almost half (47.7%) of the phishing attacks that Microsoft's detection missed were sent from compromised accounts.

Once the email lands in a recipient's inbox, the cybercriminals also hope to leverage established relationships to compel the target to take action.

PERCENTAGE BREAKDOWN OF PHISHING ATTACKS BY SENDER DOMAIN BETWEEN JANUARY – SEPTEMBER 2023 VS. 2022

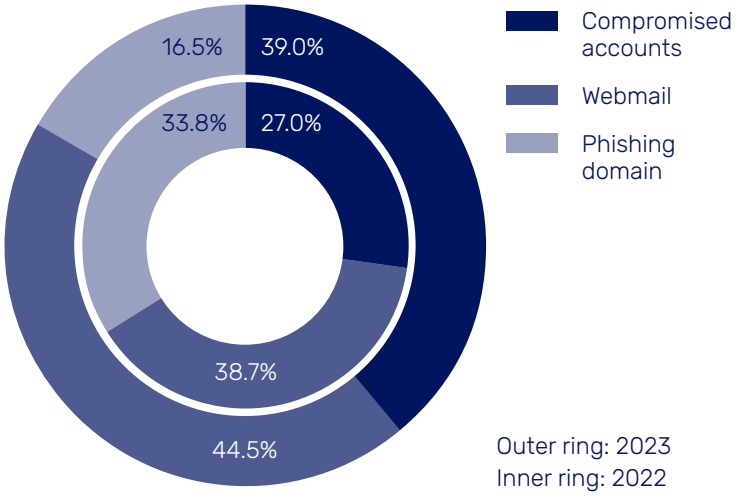

Q What is the most common type of payload?

A Links to phishing websites. Interestingly, we've observed phishing links accelerate over attachments between January – September 2023, while 'payloadless' attacks that exclusively leverage social engineering have increased by 1.1% in 2023.

PERCENTAGE BREAKDOWN OF PHISHING PAYLOADS BY TYPE BETWEEN JANUARY – SEPTEMBER 2023 VS. 2022.

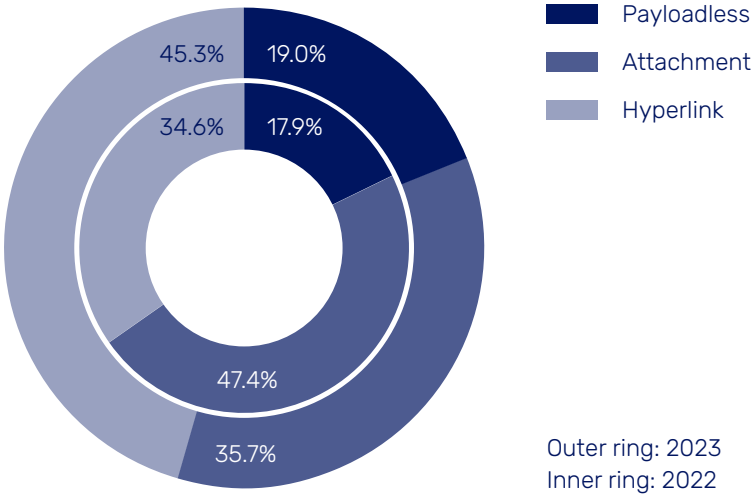

Q Which payloads bypass signature-based detection?

A All of them to different degrees. The chart shows the breakdown of primary payload contained within the emails that got through signature-based detection but were ultimately identified as phishing by Egress Defend. As phishing links are the most common payload, they were missed more frequently, followed by malicious attachments and payloadless attacks, which are seen less frequently.

THE PAYLOADS CONTAINED IN PHISHING EMAILS THAT GOT THROUGH SIGNATURE-BASED DETECTION IN MICROSOFT 365 AND SECURE EMAIL GATEWAYS

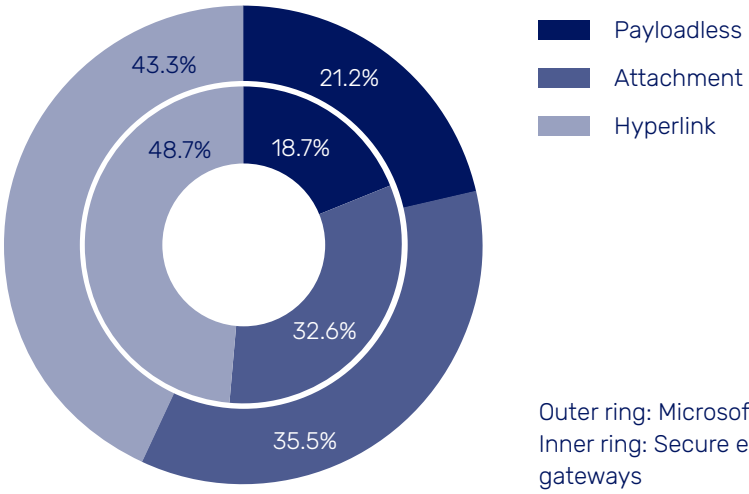

Q What are the most phished industries?

A Banking and financial services receive the most phishing emails, with legal in second place, and healthcare in third.

UNITING TO PROTECT THE HUMAN ELEMENT

As threats evolve, the cybersecurity industry needs to work together to manage human risk on email.

We hope you found this edition of the Phishing Threat Trends Report insightful. We produce it to keep you informed about changes in phishing attack trends and highlight key considerations about how to adapt your defenses in response.

At Egress, we believe that the only solution to managing human risk on email is for the cybersecurity community to work together. Only by aggregating data to derive insights can we build a holistic view of human risk that's tailored to each individual person. With that unique human risk score, we can adapt technology and training to automatically protect people when they need it most - in the moment of risk.

The Egress team would be delighted to continue the conversation about protecting your organization from advanced phishing threats.

We look forward to speaking with you soon.

GET IN TOUCH

ABOUT EGRESS DEFEND

An integrated cloud email security solution, Defend integrates seamlessly into Microsoft 365 to detect the advanced phishing attacks that get through its native security and secure email gateways. Using AI models to deliver behavior-based security, Defend detects zero-day and emerging phishing threats, attacks sent from compromised accounts, and those leverage social engineering.

Using dynamic banners applied to neutralized threats, Defend provides real-time teachable moments that continually 'nudge' employees into good security behaviours to tangibly reduce risk and augment security awareness.

Request demo

Learn more

About Egress

Egress is the only cloud email security platform to continuously assess human risk and dynamically adapt policy controls, preparing customers to defend against advanced phishing attacks and outbound data breaches before they happen. Trusted by the world's biggest brands, Egress is private equity backed with offices in London, New York, and Boston.

www.egress.com