1H CYBER ATTACK REPORT

2021

AVANAN

EXECUTIVE SUMMARY

- Our report shows that existing layers of security are far from perfect and let far too many attacks through. According to an Avanan analysis, 5% of all emails are phishing.
- Without the use of sophisticated AI, 51% of attacks would be missed by these layers and reach end-users.
- Because threats have gotten so advanced, AI is required to stop the majority of attacks missed by legacy solutions.
- Hackers are starting to target lower-hanging fruit, no longer just trying to go for the "whale".
- The Junk folder has become a key weapon for phishers.

CONTENTS

TOP TAKEAWAYS

The landscape of email attacks has changed dramatically. It is no longer a numbers game from the attackers. It is about finding the right targets and putting in a lot of effort to compromise them. Avanan has found that without the use of sophisticated AI, more than 50% of the attacks would be missed. The days of spray and pray are over. Instead, employees at all levels are being targeted.

In today's phishing landscape, impersonation and credential harvesting attacks remain king. Tactics like using non-standard characters and no or limited sender reputation remain at the top.

Hackers remain effective at finding new ways to get past default security. Employees in all sectors and at all levels need to remain vigilant.

HOW WE PRODUCED THIS REPORT

Avanan security researchers analyzed over 905 million emails.

Since Avanan works as a layer of security behind Microsoft's EOP, ATP/Defender, Google Workspace, or any SEGs, our analysis only looks at the emails these other layers did not quarantine. Our report reflects an analysis of the most sophisticated and evasive attacks in use today.

TODAY'S THREAT LANDSCAPE

When cybersecurity reporter Brian Krebs released his book, Spam Nation, in 2014, he thought that surely, within a few years, the spam and phishing problem would be significantly reduced. He thought it was possible that the problem could be completely solved.

Fast-forward seven years and phishing has not only not gone away—it's gotten more fierce. Phishing attacks remain the most widespread email threat to organizations across the globe. According to the Verizon Data Breach Investigation Report, the majority of breaches today are caused by phishing, Business Email Compromise (BEC), and credential thefts. The FBI IC3 Report found that phishing was the most common cybercrime report in 2020 and that BEC attacks have caused victims more than $1.8 billion in losses. IBM found that nearly one in five companies suffered a malicious data breach caused by lost or stolen credentials. As email has primarily moved to the cloud, a new era of phishing attacks has emerged. Now, hackers have a wider repository from which to launch devastating attacks.

Those attacks are getting more costly, too, with one analysis finding that the annual cost associated with phishing for a 10,000 user company is $3.7 million. And it's becoming more time-consuming. Research has shown that the SOC spends 22.9% of their time managing the email threat.

Further, phishing attacks have the potential to be more devastating than ever before. With the increase in cloud-based file-sharing and collaboration apps, like Dropbox and Microsoft Teams, hackers can easily compromise an email account and make their way laterally to other apps that hold sensitive data.

PHISHING VECTORS

CREDENTIAL HARVESTING (54% of phishing attacks)

The idea behind credential harvesting is for hackers to find a way to get the victim to divulge personal information. Those credentials can link to email accounts, to bank accounts or they can be Social Security or credit card numbers. These attacks work by impersonating a trusted brand or person, which then tricks the user into entering their credentials into a spoofed login page. Files will either have HTML attachments or links to websites controlled by the attacker that resembles a trusted service. After obtaining the credentials, hackers either sell the info on the black market or take over the account to access files and other sensitive information.

Credential harvesting is on the rise. Since our last report in 2019, it has risen by nearly 15%. In the previous six Verizon Data Breach Reports, credential harvesting has quickly risen to the top of what causes breaches. Even amongst malware, password dumpers (which are used to get credentials) have taken up the top spot on breach varieties, according to the report. In New Zealand, for example, credential harvesting and phishing, together, were the most commonly reported attack forms, up a whopping 76%.

And credential harvesting attempts are working. According to the 2020 Gone Phishing Tournament by Terranova, of the 20% of those likely to click on phishing links, 67.5% of those go on to enter their credentials on a phishing page. Microsoft has found that state actors are more focused on credential harvesting than in the past; they also found that in 2019 alone, more than 1 billion URLs were set up to steal credentials. Some 2 million URLs a month are being created.

Part of the reason there's so much volume is that it is effective. Take a look at the following email:

This email is a typical credential harvesting attempt. It looks convincing enough that a harried end-user will probably click. When they click on Keep Password, they are taken to this page:

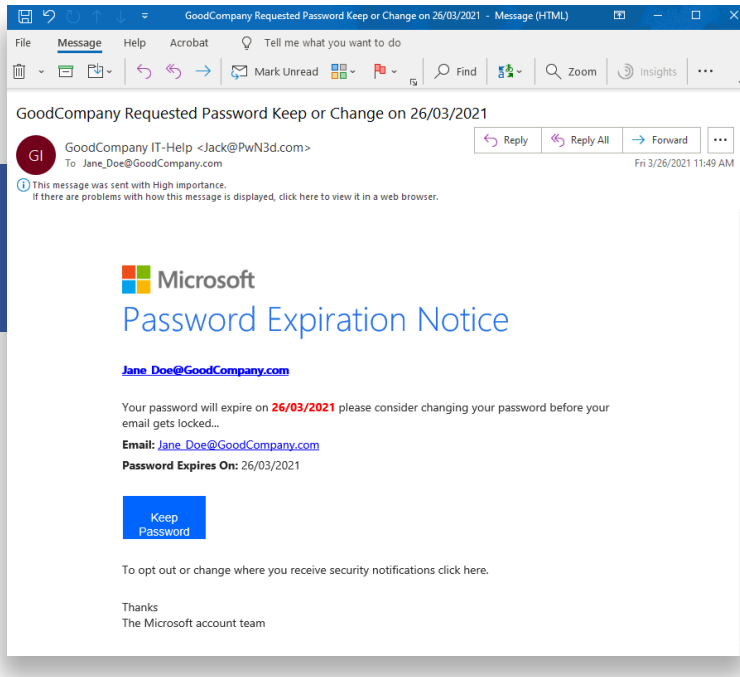

That is a traditional-looking Microsoft login page. As soon as the end-user enters their password, it's game over.

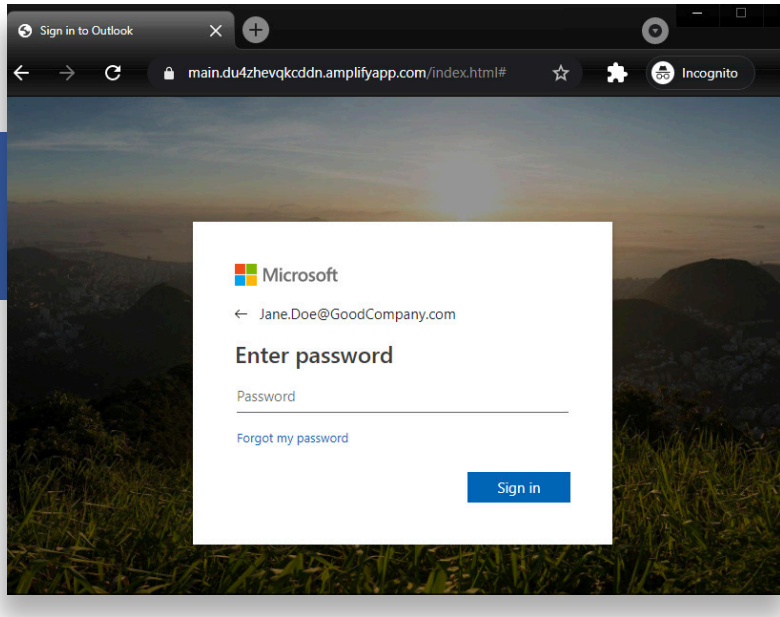

Some hackers have gotten particularly good at this. SPAM-EGY, an advanced persistent threat group that Avanan has written about extensively, dynamically renders login pages that match company logos. In a specific university-targeted campaign, the code actually imports the school's login page on the fly so that it is an exact match. All you have to do is change the URL of the email address. We tried a few well-known universities as examples:

For the average end-user, this looks legitimate and makes it easy for hackers to lure them in.

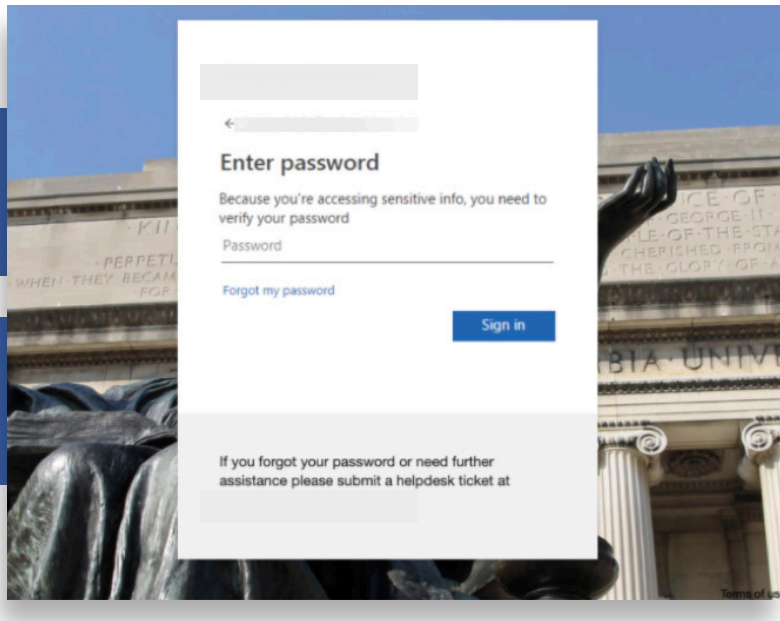

Once they do, hackers can either take over the victim's account or sell the information on the Dark Web.

Credential harvesting occurs across all attack vectors.

PHISHING VECTORS

BUSINESS EMAIL COMPROMISE (20.7% of all attacks)

Business Email Compromise (or BEC) is simple: a hacker pretends to be someone or a company that they are not. It is a particularly difficult form of impersonation to catch. Often, the emails are simple and straightforward:

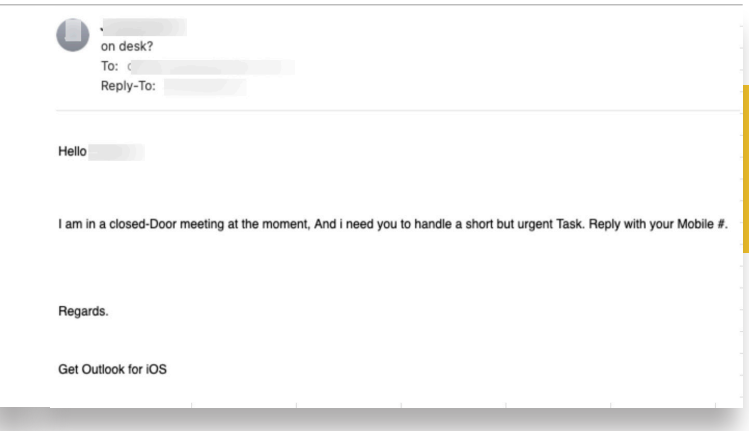

This appears as a simple request from a higher-up. It's not. It's someone impersonating an executive, hoping to get someone junior to respond.

These sorts of attacks are quickly becoming a go-to for hackers. The average BEC payment nearly doubled between the first and second quarters of 2020. It's now at $80,183, on average. Gartner found that BECs increased by nearly 100% in 2019 and through 2023, predicts that BEC attacks will continue to double each year, at a cost of over $5 billion to its victims.

Another common form is so-called Gift Card BEC scams. The Internet Crime Complaint Center tracked a 1,240% increase in 2018 of these types of attacks. And at the beginning of the COVID-19 pandemic, scammers were asking victims to buy gift cards to help purchase PPE.

BEC attacks are rising in large part because they are difficult to identify and stop. The only way to stop BECs is through internal context. Many email security solutions don't have this. For example, when a Secure Email Gateway like Mimecast or Proofpoint sees an email from the 'CEO' to the 'CFO' it will be the first time it has seen such a conversation. Even Microsoft and Google, which do have internal access, don't have the infrastructure to perform per-customer contextual analysis.

These attacks require AI to catch effectively, and Avanan's AI, combined with scanning of internal email and deployment-day analysis of one year's worth of email conversation to build a trusted reputation network, can stop BECs in their tracks.

PHISHING VECTORS

IMPERSONATION ATTACKS

Impersonation attacks are just that—the sender impersonating someone else. There are three main types: User, Domain, or Brand. Attackers will change the sender address on the email's headers to spoof the desired target. That can mean impersonating the CEO or CFO. It can mean impersonating someone from an external company. It can mean impersonating a popular and trusted brand, as well. Whatever the impersonation is, the idea is to be convincing enough that the victim gives up information or data that they would normally feel comfortable releasing.

Typically, at least according to recent thinking, hackers would try to impersonate C-level executives. And while that happens often—29.4% of all impersonation emails are in the C-Suite—what is a novel finding is that hackers have switched up tactics. Now, according to our findings, 51.9% of all impersonation emails attempted to impersonate a non-executive in the organization.

In fact, non-executives are targeted **77% more often**, according to Avanan research.

There are a few reasons behind this. One, security admins might be spending a lot of time providing extra attention to the C-Suite and hackers have adjusted. Two, non-executives still hold sensitive information and have access to financial data. There is no need to go all the way up the food chain.

This was targeted at a lower-level employee who had access to such information. When users are responding, you don't need to reach the highest levels.

Security solutions need to adjust to this new trend. Avanan protects all end-users the same, whether they are an intern or the CEO.

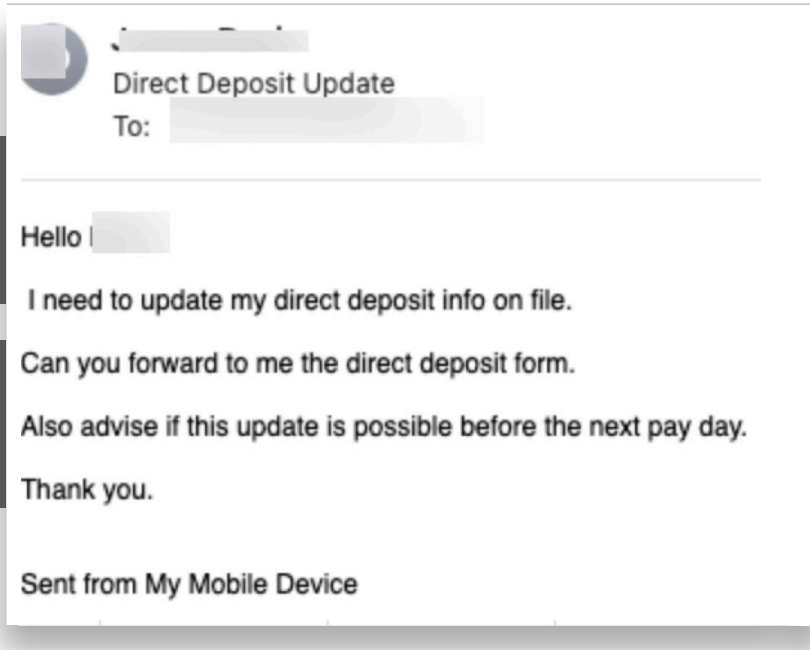

Another pervasive form of impersonation is Brand Impersonation. This is when hackers take advantage of popular brand names to fool the user. Some of the biggest brands in the world are used frequently by hackers. Consistent with other research, Microsoft takes the cake as the most impersonated brand:

Brand	Brand Impersonation %
Microsoft	43.19%
IRS	13.75%
Amazon	12.32%
Naver	10.98%
AOL	10.33%
DHL	9.43%

This is what a typical Microsoft impersonation attempt looks like:

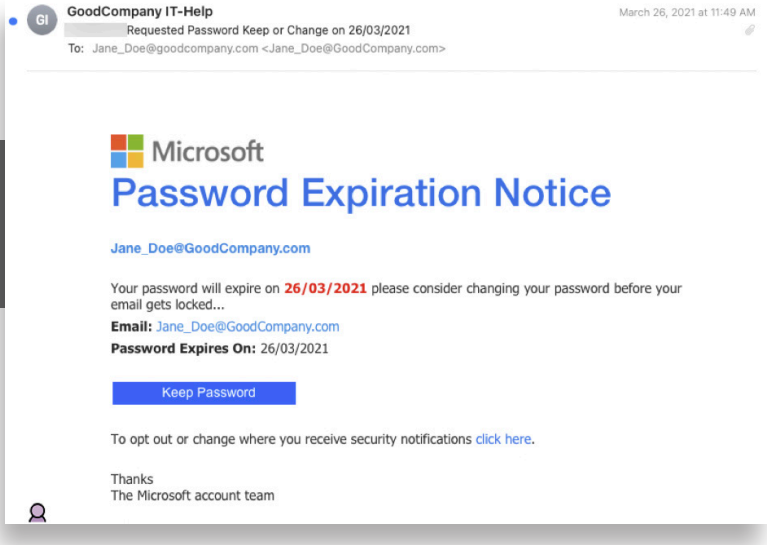

But it doesn't have to be a mega-conglomerate. This email impersonated the sunglass company Ray-Ban:

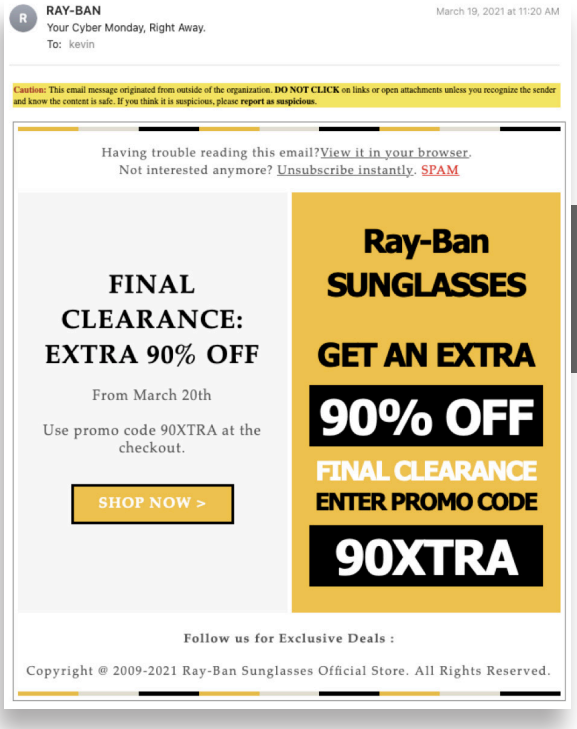

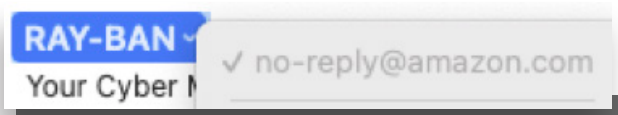

When you hover over the sender, you'll see a double spoof:

PHISHING VECTORS

EXTORTION (2.2% of phishing attacks)

Like real-world blackmail, extortion emails are usually out for money. This is a high-pressure email that usually has threats of ruining the person's life unless payment, in the form of bitcoin, is made immediately. Though the actual message content can be vague, the content is usually salacious in nature. The attack will often list something that was obtained from a data leak to add credibility.

Extortion emails can be quite graphic and quite intimidating. It's targeted to the user, and it appears like they have intimate knowledge of the user's life. That's one of the reasons it works. In 2019, complaints about extortion increased by 242%, with losses of $83 million.

In general, these emails use a spoofed domain that gets by Microsoft's natural language processing, usually by utilizing a form of obfuscation. They also typically contain a crypto-wallet address:

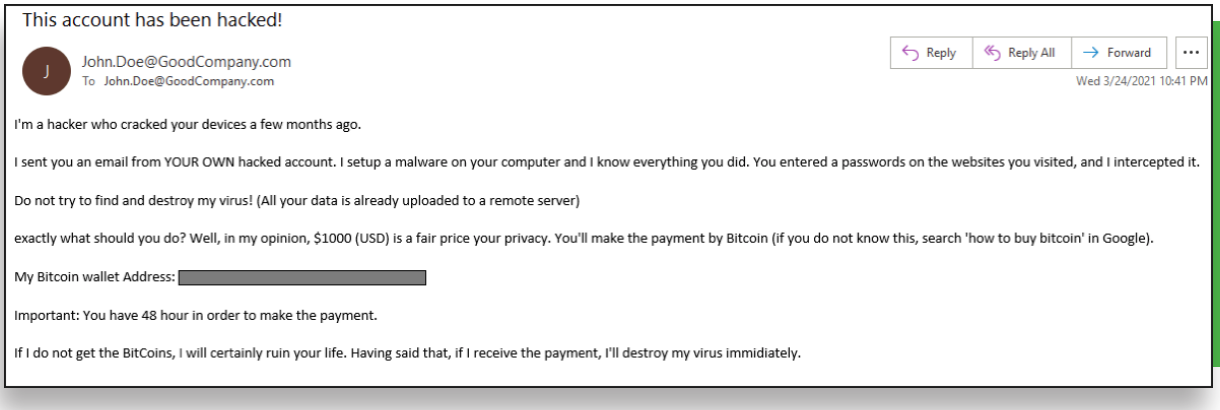

This account has been hacked!

J John.Doe@GoodCompany.com
To John.Doe@GoodCompany.com

Reply | Reply All | Forward | ...
Wed 3/24/2021 10:41 PM

I'm a hacker who cracked your devices a few months ago.

I sent you an email from YOUR OWN hacked account. I setup a malware on your computer and I know everything you did. You entered a passwords on the websites you visited, and I intercepted it.

Do not try to find and destroy my virus! (All your data is already uploaded to a remote server)

exactly what should you do? Well, in my opinion, $1000 (USD) is a fair price your privacy. You'll make the payment by Bitcoin (if you do not know this, search 'how to buy bitcoin' in Google).

My Bitcoin wallet Address: [redacted]

Important: You have 48 hour in order to make the payment.

If I do not get the BitCoins, I will certainly ruin your life. Having said that, if I receive the payment, I'll destroy my virus immidiately.

Of those who received an extortion email, extortion accounted for 38.04% of all phishing emails.

INDUSTRY-BASED PHISHING

The most attacked industries, according to our research, are IT, healthcare and manufacturing/ industrial. IT saw by far the most, with over 9,000 phishing emails in a one month span, out of an average of 376,914 total emails. Healthcare saw over 6,000 phishing emails out of an average of 451,792 total emails; manufacturing saw just under 6,000 phishing emails out of an average of 331,184 total emails Why are these industries the most targeted? Where there is rich data, there will be attacks. These industries, in particular, hold incredibly valuable data, from health records to social security numbers and more. Combined with that is the fact that healthcare and manufacturing, in particular, tend to use outdated tech and often have non-technical board of directors. In healthcare, in particular, the industry is largely unprepared. According to a study, 87% of organizations say they don't have the proper personnel to defend against attack; another study found that 32% of hospital personnel haven't received the proper security training. Though every industry gets attacked, the ones that hold the most data are the most at risk.

All industries are attacked. Even those with lower numbers, like non-profit, are still being hit. The vast range of attacks show how, in general, attackers are changing up their targets. Regardless of sector, hackers are hitting users who they have found are more susceptible to phishing attacks. It represents a broader trend we've noticed in phishing, whereby attackers will target those who are perhaps more susceptible to such attacks, instead of those whose defenses are incredibly strong.

JUNK MAIL ANALYSIS

In May of 2020, Avanan published a blog about the trend of "Dumpster Diving" in the junk folder. Because Microsoft flags an unusually high percentage of legitimate emails as phishing, many organizations have decided to send detections to the Junk folder. The "Dumpster Dive" policy sees marketing emails, subscriptions, and phishing attacks commingled in the Junk folder and easily accessible to end-users. For Microsoft, SCL scores of 5,6, and 9 will be sent to a user's Junk folder. Once there, everything—good and bad—is up for grabs. As one CIO of a Fortune 500 company told us, "You now have monthly subscriptions, newsletters, and targeted phishing attacks in your spam folder, and you have to leave it up to the end-user to decide which ones are safe to open."

The numbers bear this out and though it also occurs with Google, Microsoft puts **89% more emails in Junk than Google does**.

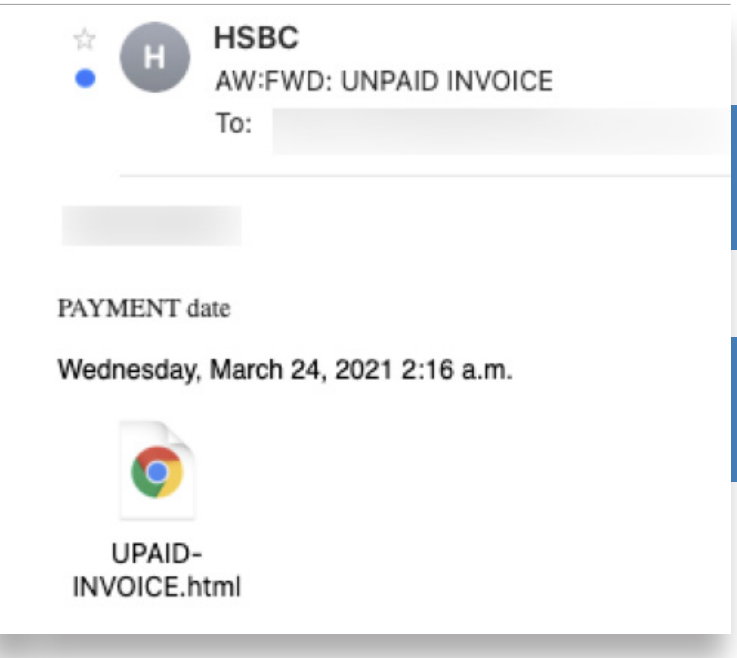

Microsoft:

4.11% of all emails in an EOP junk folder is Phishing
63.5% of Phishing emails are in the junk folder

Google:

2.16% of all emails in a Google Workspace Junk folder is phishing
67% of Phishing emails are in the Junk folder

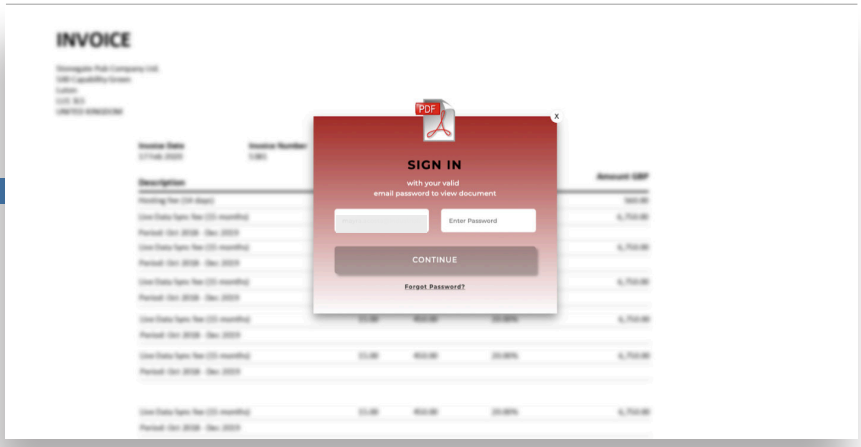

Allow lists

One of the more intriguing findings from our analysis is the role that misconfiguration plays in phishing. We found that **8.14%** of phishing emails ended up in the user's inbox simply because of an allow or block list misconfiguration. This is an increase of **5.3%** from the 2019 Global Phish Report.

The problem gets worse depending on the security solution in use. When sitting behind an SEG, we found that **15.4%** of email attacks are on an Allow List.

With Avanan, there are no prerequisites or complicated rules to follow. Just connect Avanan and it begins working. This process is faster with mature AI platforms and Avanan's AI has a unique advantage because Avanan is trained with a dataset of emails that are missed by the email layers that come before Avanan.

In the above example, "id.apple.com" was in this customer's allow-list.

LINK ANALYSIS

Avanan found that non-standard characters are used in 50.6% of phishing links, by far the most commonly used:

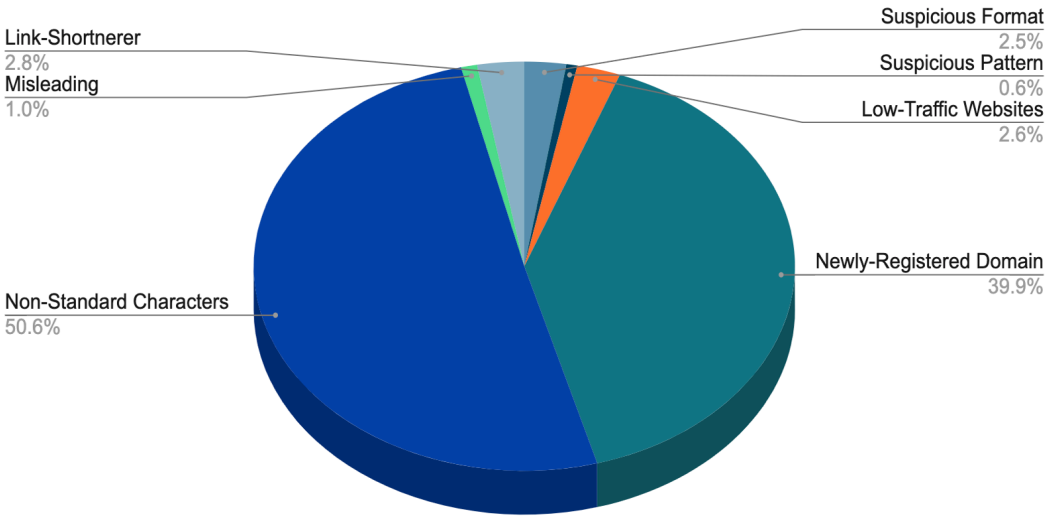

These non-standard links are used to bypass link scanners in Microsoft and SEGs.

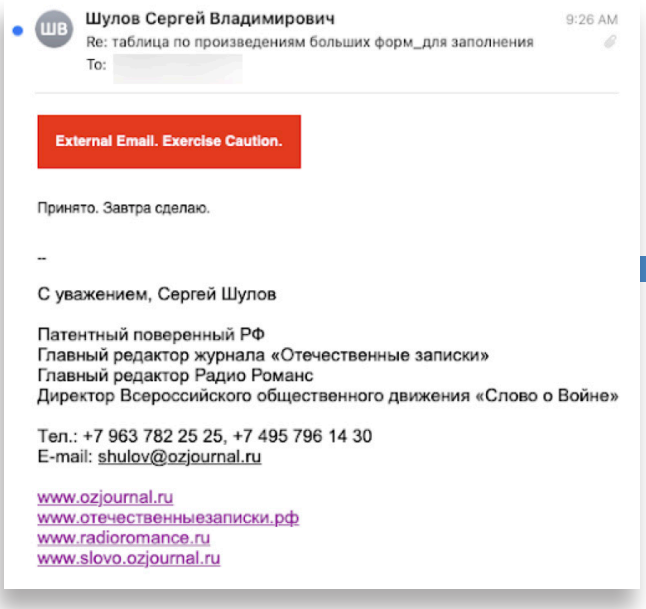
Шулов Сергей Владимирович 9:26 AM
Re: таблица по произведениям больших форм_для заполнения
To:

External Email. Exercise Caution.

Принято. Завтра сделаю.

--

С уважением, Сергей Шулов

Патентный поверенный РФ
Главный редактор журнала «Отечественные записки»
Главный редактор Радио Романс
Директор Всероссийского общественного движения «Слово о Войне»

Тел.: +7 963 782 25 25, +7 495 796 14 30
E-mail: shulov@ozjournal.ru

www.ozjournal.ru
www.отечественныезаписки.рф
www.radioromance.ru
www.slovo.ozjournal.ru

For Avanan, the usage of a non-standard character is a red flag. Any non-standard character in a link increases Avanan's AI score and when combined with other AI indications, it's an effective way to determine if a link is malicious.

SENDER REPUTATION

An easy way to determine if an email is suspicious is by looking at sender reputation. It's no wonder, then, that **84.3%** of all phishing emails do not have a significant historical reputation with the victim. Further, **43.35%** of all phishing emails come from domains with very low traffic.

Low Reputation:

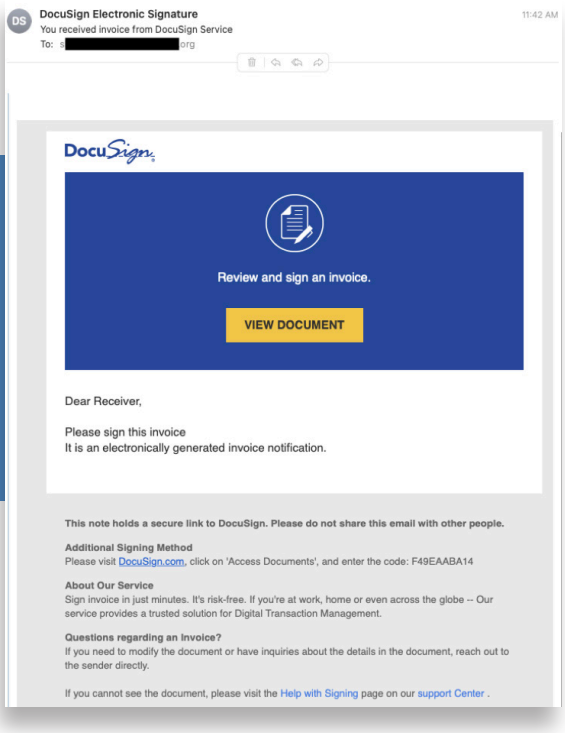

High Reputation:

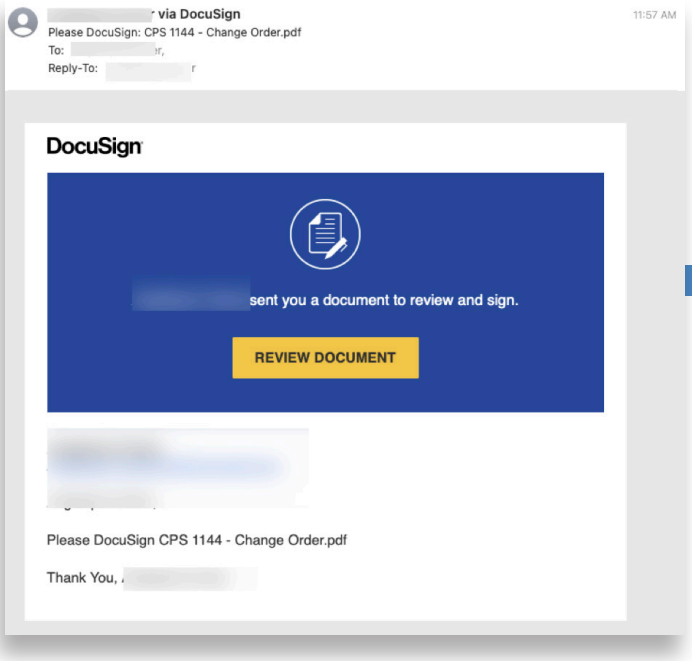

High reputation and low reputation emails look identical. End-users are tasked with figuring out which is real and which is not. Therefore, it's imperative that all security vendors use Sender Reputation as one of the key metrics when identifying phishing.

THE FUTURE

Phishing is evolving every year. Hackers will continue to find vulnerabilities in existing security systems. They will continue to find new and innovative ways to bypass this protection and get into the inbox. And once they get into the inbox, they are finding unique and sophisticated ways to get the user to click and enter the information they're looking for. Hackers will continue to focus on industries that don't have proper defenses, preferring to gather up the lowest hanging fruit at a higher rate than hoping to bring in the big whale.

On top of this, new factors are at play. There will be an increase in collaboration software as attack vectors. Apps like Microsoft Teams and Slack grew exponentially during the pandemic and became a lifeline for many companies as they navigated a remote-work world. However, both of these apps are inherently insecure and hackers are beginning to target these apps for the often valuable data they share. In a survey of IT leaders that Avanan conducted at the end of 2020, 76.1% of respondents strongly agreed or agreed that vulnerabilities in both platforms present a security risk and that they would have to adopt mitigation technology within the year.

On top of that, email isn't going away as an attack vector. So long as phishing works, hackers will target it. Instead of just protecting emails, companies now have to protect all the apps in which they do business, from collaboration to file sharing and everything else where data is shared and stored.

One way to combat this is with phishing training. It's another layer of defense, but it is not a panacea. One study found that phishing training awareness wears off, and a new round of education is needed twice a year.

That's why training needs to be paired with advanced AI. **Avanan's analysis found that 50% of all attacks would be missed without AI**. As email continues to migrate to the cloud, and as the market continues to shift away from SEGs towards API security, leveraging that positioning by training AI on the most sophisticated attacks is the best way forward.

THE AVANAN DIFFERENCE

Avanan differentiates itself with its AI detection methodology. Avanan implements over 300 (and growing) Indicators of Attack (IoA) to determine whether an email is phishing. A non-exhaustive list of some of the indicators Avanan uses:

- Phishing language in an email's subject and body
- Encoded content, such as scripts to encode or decode Base64, Morse, etc
- HTML obfuscation methods, such as ZeroFont and baseStriker
- Existence of a Crypto wallet address

Additionally, Avanan's positioning, behind your other layers of security, means it is trained on phishing emails that got past Microsoft EOP or ATP, Google, and any other SEG a client may have installed. By doing so, Avanan is better equipped to stop the newest, most sophisticated, and evasive attacks in the wild today.

This patented approach has been a boon for customers. Why? Because Avanan is highly accurate and stops sophisticated attacks, all while saving your team time from managing the email threat. Avanan customers see a 99.2% reduction in phishing attacks reaching the inbox. The SOC is less busy, with a 71% reduction in phishing-related alerts sent to that department.

Additionally, in a study of 360 million emails, Avanan was 15x more effective than legacy gateways like Proofpoint, Mimecast, and Barracuda, and 18x more effective than Microsoft ATP.

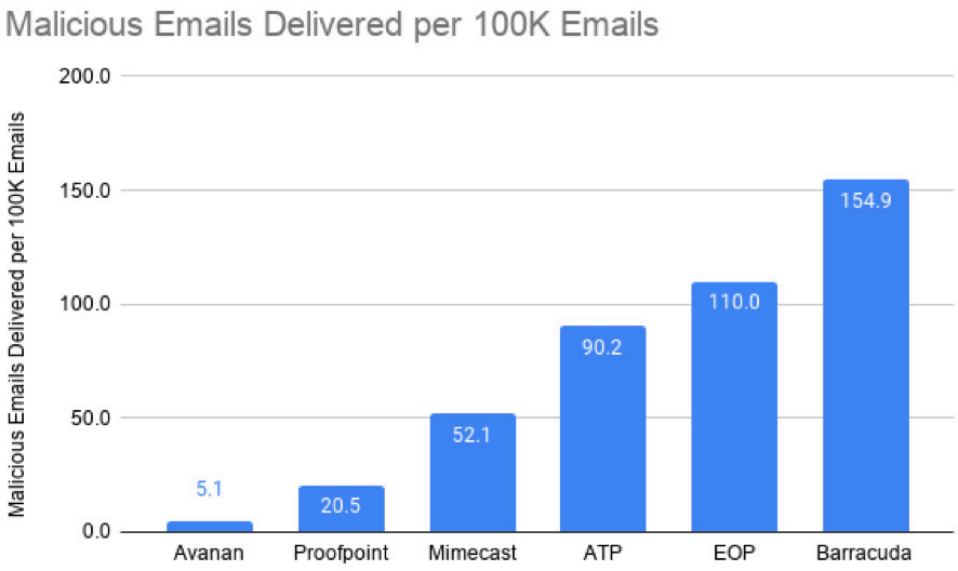

The state of global phishing is precarious. Threats are more sophisticated than ever. Payments are higher. New attack trends are evasive and effective. New attack vectors are gaining prevalence.

All of this has the potential to make life for corporations and their employees miserable. However, with proper protection like Avanan, CEOs and CISOs can properly confront phishing from the boardroom to the inbox, with industrial-strength effectiveness.

ABOUT AVANAN

Avanan catches the advanced attacks that evade default and advanced security tools. Its invisible, multi-layer security enables full-suite protection for cloud collaboration solutions such as Office 365™, G-Suite™, and Slack™. The platform deploys in one click via API to prevent Business Email Compromise and block phishing, malware, data leakage, account takeover, and shadow IT across the enterprise. Avanan replaces the need for multiple tools to secure the entire cloud collaboration suite, with a patented solution that goes far beyond any other Cloud Email Security Supplement.