



Annual Threat Report 2024

Cybersecurity insights,
resilience recommendations,
and predictions



“We’re in this together and strong, collaborative information sharing is key.”



Dave Merkel
CEO, Expel

Letter from the **CEO**

Expel’s operators do a massive amount of analysis, triage, and complicated problem-solving—stopping intricate attacks every single day.

It’s our responsibility as operators to take lessons learned from our varied customer base and share those insights with others in our community in the interest of banding together against our common adversaries.

For each of the top trends we explore, we’ll tell you what we saw, how to detect these threats and protect your organization, and what to watch for in 2024. We cap it off with some thoughts and predictions for cybersecurity this year, courtesy of our leadership team and some of the smartest folks in the industry (who also happen to be Expletives).

We delve into the data to identify trends and patterns in the security events we’re seeing and translate them into strategic guidelines your organization can put into play right away. Whether it’s hardening your environment against specific known threats or improving measures to detect bad actors who find their way past the first lines of defense, we’re in this together and strong, collaborative information sharing is key.

A handwritten signature in black ink that reads "Dave Merkel". The signature is stylized and cursive.

“Effective, modern SOC operations enable and empower analysts to make effective real-time decisions to protect the enterprises entrusted to them.”



Daniel Clayton
VP, Security Operations

Letter from the VP, Security Operations

Thanks for taking the time to review our Annual Threat Report. I know firsthand that the day of a cyber operator is a busy one and it's not easy to find time to read and digest such an extensive report.

The challenges you face every day are complex, fast-moving and constantly evolving, so it's our goal to provide a resource packed with intel that is both valuable and actionable for protecting your organization.

Much of the information included in this report originates from our technology stack, but it's the experience and the expertise that our analysts bring to the fight which enables our SOC team to go beyond the technology and protect our customers. These people are the best in the business, and they work behind the scenes to keep our customers (and Expel) safe.

People are at the core of the work we do, and although a modern security strategy and the right technology are critical, without a high-performing SOC, security teams will lose more often than they win. Effective, modern SOC operations enable and empower analysts to make effective real-time decisions to protect the enterprises entrusted to them. We believe this passionately and it's for this reason that we're sharing this information and intelligence with you, the cybersecurity community.

A handwritten signature in black ink, appearing to read 'D. Clayton', written in a cursive style.

Contents

- Executive summary..... **2**
- 2023 by the numbers **6**
- Identity threats..... **7**
- Cloud platform threats..... **15**
- Computer-based threats **19**
- Phishing threats..... **27**
- Annual spotlight **30**
- Looking ahead to 2024... **31**
- 2023 at Expel **33**
- Reference highlights..... **34**



Executive summary

Key insights and takeaways

Now in its third installment, the trends, predictions, and recommendations discussed in this report are based primarily on incidents our security operations center (SOC) identified through investigations into alerts, email submissions, vulnerability disclosures, and threat-hunting leads spanning January 1 to December 31, 2023. We relied on a combination of time-series analyses, statistics, customer input, and analyst expertise to generate these insights.

To make it easy to follow along, we've broken this report into four different threat types: identity, cloud, computer-based, and phishing. Keep in mind, though, that these threats are tightly related. Compromise of user identities can occur through phishing and can impact an organization's cloud assets and endpoints, or malicious activity on an endpoint may be indicative of an exploited vulnerability. We'll highlight that interconnectivity throughout.

OUR TOP TAKEAWAYS:

Identity-based incidents dominate three years in a row

Identity threats accounted for **64%** of all incidents our SOC investigated and increased in volume by **144%** from 2022. Of those incidents, **60%** were unauthorized email logins and **40%** were authentications to identity platforms, like Microsoft Entra ID (formerly Azure Active Directory), Okta, Ping, and Duo.

Of the orgs targeted with an identity attack:

- **35%** experienced more than one incident (up from **24%** in 2022)
- Organizations saw an average of **eight** identity-based incidents over the year
- One nonprofit organization was targeted **255 times** (up from **104 times** in 2022 at the same organization)

Identity threats accounted for 64% of all incidents our SOC investigated and increased in volume by 144% from 2022.

Sixty-nine percent of identity-based incidents involved malicious logins from suspicious infrastructure, which are hosting providers or proxies that aren't expected for a user or organization. We've noted in past years a shift toward using more proxies, VPNs, and hosting providers to bypass network zone and conditional access policies, and we expect we'll continue to see this trend until organizations consistently put effective roadblocks in place—such as phish-resistant multifactor authentication (MFA) and policies to block suspicious logins.

Phishing-as-a-service (PhaaS) drives identity-based incidents

The increased volume of identity incidents noted above is a direct result of more phishing platforms becoming available on the dark market. Several of these harvesters can both pre-fill the intended victim's email address and load the appropriate branding and background for the target organization's login page, [making these pages look convincingly like the expected login page](#).

It's not surprising to see another branch of cybercrime-as-a-service (CaaS) gain traction—it's accessible, convenient, and available with a low barrier to entry. Unfortunately, that means it's not going away anytime soon and we're only likely to see the attack vector grow in 2024.

The rise of QR code phishing

The [recent uptick in QR code phishing](#)—or qishing—is cause for concern because it takes the activity off endpoints and puts it on mobile devices, which may not be managed with tight security controls—allowing the attacker to bypass endpoint security. We found that Microsoft’s threat intelligence can reliably identify phishing pages created with PhaaS platforms. However, Microsoft’s effectiveness in identifying these pages may have inadvertently contributed to an increase in threat actors’ use of QR codes for phishing.

Personal devices typically lack the same level of protection as endpoint devices, meaning threat actors may be able to bypass the usual security barriers more easily.

Personal devices typically lack the same level of protection as endpoint devices, meaning threat actors may be able to bypass the usual security barriers more easily. It’s pivotal users become as wary of unexpected QR codes as they are URLs (or at least, as they should be).

Threat groups favor social engineering tactics to mimic employees and target organizations

“[The Com](#),” a hacking group including the actors tracked as “Scattered Spider,” was responsible for the largest number of targeted identity attacks our SOC investigated this year. These attackers primarily targeted Okta and Microsoft accounts, attempting to abuse password reset mechanisms and MFA push fatigue to gain access.

We observed two main attack tactics:

- Calling into an organization’s helpdesk
- Abusing self-service password reset mechanisms

When calling into the helpdesk, the actors impersonate staff and request that their passwords be reset. If either the helpdesk or self-service password request attempts are successful, the threat actor sends MFA pushes to the real user. If the user accepts the MFA push, the attacker gains access to the account.

Attackers spend a stunning amount of time conducting [thorough research on possible victims](#) to mimic their speech patterns. Of note, “The Com” primarily targeted customers in the financial services industry in the second half of the year. We expect this attack style to continue into 2024, especially as generative AI has the potential to make this kind of impersonation even easier.

Cloud infrastructure incidents trend up, with stolen or leaked cloud credentials (aka secret) exposure as the biggest and most frequent risk

This year, we noted a **72%** increase in cloud infrastructure incidents, roughly consistent with what we saw in 2022 and continuing the upward trend since we began supporting Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Kubernetes. Exposed credentials (or secrets) were the leading root cause of cloud infrastructure incidents (**42%**). Publicly exposed or stolen credentials allow attackers to maintain persistent access to the cloud environment with the permissions tied to that identity or role.

Common misconfiguration of Amazon Cognito (aka AWS Cognito) allows attackers to gain direct access to AWS control plane

Amazon Cognito (also known as AWS Cognito) enables user sign-up and authentication to apps, but administrators frequently misconfigure the service. In these cases, an attacker can do things like create new accounts with excessive permissions or even access AWS credentials associated with the Cognito Identity Pool.

This year, security researchers released a few [Amazon Cognito auditing tools](#). While these help operators identify misconfiguration, threat actors can also take advantage: we’re seeing more of these AWS Cognito focused attacks seemingly enabled by the auditing tools. But the misconfigurations that lead to these attacks are common—and thus, effective—for a reason. We wrote more about what to know and understand about Amazon Cognito in this [recent blog post](#).

More than half of all malware incidents presented an immediate, significant risk

We separate malware into categories based on its functionality. In more than half the malware incidents, the malware deployed presented an immediate and significant risk to the environment (what we would classify as “high-risk”)—including risks of pre-ransomware activity or exfiltration. Pre-ransomware accounted for **57%** of the malware incidents our SOC investigated.

In more than half the malware incidents, the malware deployed presented an immediate and significant risk to the environment.

The most frequent malware cases we classified as pre-ransomware included:

- Gootloader—**17%** of high-risk malware incidents
- Qakbot—**12%**
- SocGhosh—**11%**

It’s worth noting that these were also the top pre-ransomware threats we reported in both 2021 and 2022. The skilled actors behind these threats have been active for a while; they’ll likely continue to be active for the foreseeable future.

Attackers trick users with sneaky infostealers downloaded from malicious ads with cloned versions of legitimate websites

Infostealers—malware that accesses sensitive data, like passwords, cryptocurrency wallets, and other information stored in the user’s browsers or on their devices—provide a wide range of options for an attacker.

While infostealing malware can be targeted, threat actors most frequently deployed it opportunistically. Adversaries often take advantage of a user looking to download software. “Cracked” software, or software that generally requires a license but attackers have modified to circumnavigate that requirement, represents the highest risk. In these situations, users are more likely to act against their better judgment in downloading software from a sketchy website to avoid buying a paid version.

This year, we saw a lot of infostealers downloaded from **malicious ads**. These ads appear at the top of search results and often imitate productivity or IT software. Clicking the ad directs the user to a cloned version of the legitimate software’s website. When the software downloads and runs, the malware executes the malicious payload—sometimes even installing the legitimate software to avoid suspicion.

Shadow IT represents a growing concern for security teams

We frequently saw attackers leveraging both search engine ads and SEO poisoning to guide users to download malicious payloads. These attacks target both Windows and MacOS systems, and the main types of malware used in these attacks are remote access tools (RATs) and infostealers.

Common “lures” include:

- Software used by IT teams (such as Advanced IP Scanner, Nmap, WinSCP)
- Productivity software (such as Notion, Notepad++, and PDF manipulation tools)

Shadow IT attacks became a fad for threat actors in 2023. Google and Microsoft actively work to combat the tactic, but we expect malvertising to remain popular into 2024 (and, ultimately, for as long as it works).

Zero-days emphasize importance of defense-in-depth and strong layers of control

Here are the top vulnerabilities our SOC identified based on incident frequency:

- Progress’ MOVEit Transfer (CVE-2023-34362)
- Adobe ColdFusion (CVE-2023-26360)
- Citrix NetScaler ADC, also known as CitrixBleed (CVE-2023-4966)

Properly configured and implemented security controls—or the lack thereof—directly affect the impact of these vulnerabilities on an org. All of these vulnerabilities were zero-days, but there’s good news: security teams can identify zero-day exploitation when it happens thanks to defense-in-depth and by understanding common attacker tactics.

Bad actors trended toward script-based files for pre-ransomware initial access

JavaScript made up the highest volume of files at **39%**, but we also saw lots of other scripting types, including executable (EXE--**20%**), LNK (**12%**), and VBS (**7%**) files.

JS and VBS are text files that receive less scrutiny from automated analysis mechanisms than EXE files. Unlike EXE, attackers can hide their functionality among benign content or by command obfuscation. By default, the native Windows program, wscript.exe, will execute these scripts when a user double-clicks the files. Attackers use this setting to their advantage, so we highly recommend changing this default setting, which can be done using a group policy object (GPO).

Witness the power of secure-by-default to fully stop threat vectors

At the very end of 2022, we saw the beginning of [a new trend toward leveraging OneNote files](#) in attacks. In January 2023, Microsoft pushed a patch to slow down the exploitation, and ultimately implemented much stronger controls in March, virtually killing the technique.

This not only illustrates how Microsoft can eliminate a threat vector, but it also proves that if organizations adopt tight controls, they also have the power to fully stop a threat vector.

Hospitality, technology, and financial services stand out as heavily targeted industries

Hospitality held onto the top spot for industries targeted by phishing attacks for the second year in a row with the highest volume at **55%**. It's followed—but not closely—by travel (**12%**), technology (**9%**), financial services, and healthcare (**5%** each).

Hospitality held onto the top spot for industries targeted by phishing attacks for the second year in a row with the highest volume of targeted attacks at 55%.

Hospitality, technology, and financial services also made the list of top industries where we identified the most high-risk malware and identity incidents. We recommend that organizations in heavily targeted industries take extra security measures, including phishing training, strong endpoint defenses, and regular reviews of the environment based on user-identified malicious emails.

Infostealer campaign targeting hospitality customers since 2022

Since 2022, our SOC has observed an infostealing campaign targeting our hospitality customers with the goal of gaining administrative access to sites like Booking.com to commit fraud.

The campaign has developed over time, but it generally looks like this: an attacker uses a Gmail account to request information about a booking, ask for help, or to lodge a complaint. Instead of an attachment, the email contains a link to a file storage service such as Dropbox, Google Drive, and Mega.nz.

In most situations, the files in the storage service are a file archive (ZIP, RAR, etc.), which is password-protected to prevent the storage provider from scanning its contents. When the user opens the file, the archive typically contains an inflated EXE file **which**, is an infostealer.

Though this campaign targets hospitality companies, understanding it is extremely valuable for organizations in other industries, too. As a common threat vector, infostealing malware can affect any organization.




2023 by the numbers

Incident types detected by the Expel SOC

When we examine the breakdown of incident types in 2023 vs. 2022 (Chart 1), here are some interesting observations:

- Identity continues to reign supreme, with a 14% increase from the previous year. Our analysts work multiple identity-based incidents a day and we expect this trend to continue upward.
- Of note: while malware as a percentage of overall incidents decreased by 25% in 2023, the potential impact of both high-risk and latent-risk malware should not be discounted.
- Phishing incidents tripled from 2% in 2022 to 6% in 2023.
- The percentage of authorized penetration tests and red teams we investigated decreased 43%. While we can only speculate, this could indicate that customers are dealing with tighter budgets and not performing these exercises as often.


Chart 1: Breakdown of incidents detected by the Expel SOC in 2022 and 2023



In Chart 2, we look at the incidents detected by technology type in 2023 compared to 2022. Some takeaways include:

- Year over year, we see attackers targeting cloud structure, as cloud incidents increased by 12% in 2023. Remember that monitoring your cloud is like planting a tree; the best time to start was five years ago. The second best time is right now.
- Similar to malware in the chart above, endpoint incidents fell as a percentage of overall incidents in 2023.
- We observed a slight decrease in initial alerts coming from network and SIEM as a percentage in 2023.

Chart 2: Initial alert source





Identity threats

Targeting email and identity platforms


Identity-based incidents were the most common attacks we saw across our customer base in 2023, accounting for 64% of all incidents handled by our SOC. The total of identity-based incidents rose 144% compared to 2022.

Chart 1 shows that within these identity incidents, 60% were unauthorized email logins (pre-BEC), and 40% were authentications to identity platforms (BAC), like Microsoft Entra ID (formerly Azure Active Directory), Okta, Ping, and Duo. While email remains the top target for attackers, we urge security teams to consider attackers that target identity platforms a serious threat, too.

2023's top threat

Over the course of 2023, the Expel SOC saw an overall increase in the number of identity incidents as the year progressed, until we saw a dip in frequency at the end of the year (perhaps due to normal seasonality). Take a look at [Chart 3](#) on page 9 to see the peaks and valleys, as well as the overall increasing trend.


Chart 3: BEC (email compromise) vs. BAC (application compromise) in 2023



WHAT WE SAW IN OUR PREVIOUS ANNUAL REPORT

- Representing 50% of all incidents, BEC attempts were the top threat to our customers in 2022.
- 60% of attempts to compromise cloud identity provider credentials (aka BAC activity) happened in Okta; 17% of the remainder were in Box and 6% were in Ping and OneLogin.
- Nonprofit, retail, and entertainment were the industries most targeted by BEC attacks.
- Of the organizations targeted by BEC threat actors:
 - 53% experienced at least one BEC attempt.
 - 24% experienced at least three BEC attempts.
 - One organization was targeted 104 times in 2022.
- BEC attackers moved away from authenticating via legacy protocols to bypass MFA in Microsoft 365. Instead, they increasingly adopted the use of frameworks such as Evilginx2 to steal login credentials and session cookies.
- BEC threat actors successfully gaining access to M365 accounts trended up. A main culprit was adversary-in-the-middle (AiTM) phishing to steal session cookies.

Chart 4: The increase in identity incidents throughout 2023



The increased volume of identity incidents is a direct result of more phishing platforms becoming available on the dark market. “Phishing-as-a-service (PhaaS)” platforms allow a buyer to easily deploy convincing credential harvesters for a phishing campaign. Several of these harvesters can both pre-fill the intended victim’s email address and load the appropriate branding and background for the target organization’s login page, [making them look convincingly like the expected login](#) page.

“Phishing-as-a-service (PhaaS)” platforms allow a buyer to easily deploy convincing credential harvesters for a phishing campaign.

Risks of identity compromise

It goes without saying that identity compromise is serious. According to the Federal Bureau of Investigation’s (FBI’s) most recent Internet Crime Complaint Center (IC3) statistics, Americans reported a combined loss of [\\$2.7 billion over 2022](#) due to BEC attacks that weren’t stopped in their early stages. BAC was the attack type behind the much-publicized [compromises at Uber, Nvidia, and Rockstar Games](#); in those incidents, attackers compromised user identities to gain access to sensitive data, sell it, and use it for extortion.


At Expel, mitigating identity threats is a priority. We understand many of these attacks attempt to leverage access as soon as possible, and we pride ourselves on our

ability to detect and respond quickly to secure an account. Our goal is to respond to the compromise in real time as much as possible to prevent the attack from playing out. However, by taking action and stopping an attack, it becomes impossible to understand the full intent of that attack. In most situations, we classify the attack based on the type of account compromised.

As such, we classify evidence of a compromised user password as an identity incident. In some instances, MFA or conditional access blocks logins; however, we treat these with the same seriousness as an incident due to the high risk that would exist were the incident to play out. Many authentication blocking methods are but a small speed bump for attackers and require our analysts to take further action to prevent potentially successful attempts. In fact, we’ve seen many instances when an attacker tried to login, was blocked by conditional access based on geolocation or MFA, and immediately switched to a bypass method, like a VPN or legacy protocol, resulting in successful login.

Furthermore, compromise of credentials is evidence of another compromise: users may have entered credentials into a credential harvester, had their credentials stolen by an infostealer (more on these later), or been exposed due to a weak password or a previous data leak. Organizations should thoroughly investigate any situation where users could have unknowingly compromised their passwords.

Chart 5: Top 10 targeted industries



Identity incidents by industry verticals

Throughout 2023, threat actors targeted non-profit organizations most, followed by technology companies and financial services.

Chart 4 shows the percentage of identity incidents our SOC saw, broken down by industry. Like in past years, this breakdown shows that threat actors don't discriminate—while some industries are targeted more often than others, none are immune.

Identity targeting frequency

Identity incidents affected 58% of our customers (up from 53% in 2022). But if we dig into the data and look at the frequency our customers are targeted, we get a clearer picture of the pervasiveness of these attacks.

Of the organizations targeted with an identity attack:

- 35% experienced more than one incident (up from 24% in 2022).
- Organizations saw an average of eight identity-based incidents over the year.
- One non-profit organization was targeted 255 times (up from 104 times in 2022)

One non-profit organization was targeted 255 times (up from 104 times in 2022)

This data illustrates that some organizations are targeted more than others, and those with a lower public profile and more compensating controls are less likely to experience identity attacks.

Identity incidents by source technology

We identified 88% of incidents based on signal from Microsoft products. This high targeting rate is due to Microsoft's large business email and identity market share, which makes the company's identity products an attractive target—similar to the way Windows operating systems are highly targeted by malware because of Microsoft's prominent place in end-user operating systems.

Okta was a distant second place to Microsoft, being the source of 9% of identity incidents. Google Workspace (recently rebranded from G Suite), Duo, OneLogin, Ping, and a few others make up the remaining technologies.

Malicious login sources

Sixty-nine percent of incidents in 2023 involved malicious logins from suspicious infrastructure, which are hosting providers or proxies that aren't expected for a user or organization. These logins may not be geographically suspicious but are abnormal for the specific user. When we analyze abnormal logins, we're looking at the details of the IP address associated with the login against the user's or account's baseline behavior.

Just over 10% of malicious logins came from VPNs, and the rest (about 21%) were identified through geolocation alerts. See [Chart 5](#) for the breakdown.

This is consistent with a trend we noted in 2022—namely, that attackers were shifting to using more proxies, VPNs, and hosting providers to bypass network zone and conditional access policies. The takeaway? Organizations need to take any abnormal logins seriously.

Chart 6: Breakdown of malicious login sources




Image 1: Expel Workbench™ IP lookup results for an anomalous login

Outcome
What do we know about IP 196.247.229.187?

- 196.247.229.187 is associated with the following Organizations: **PACKETEXCHANGE`Packet Exchange Limited`fiber grid inc**
- Is 196.247.229.187 a TOR node?: **False**
- Is 196.247.229.187 a VPN server?: **False** | [Cross reference with Spur](#)
- Is 196.247.229.187 a Proxy server?: **False**
- Is 196.247.229.187 a hosting provider?: **True**
- 196.247.229.187 originates from **Tallinn, Harjumaa, 10111 EE**
- 196.247.229.187 is associated with the following ASN **58065**
- iphub classifies the ip as **Benign**
- greynoise has tagged 196.247.229.187 for the following reason: **The IP has never been observed scanning the Internet.**

Let's look at an example: here's a user logging in from the IP address 196.247.229.187, associated with Packet Exchange, which differs from the expected baseline.

We compare this information with the user's baseline and often we're able to identify clearly if the origin is normal or abnormal for the user.

Image 2: Results from automated review of user behavior

[-----| Historical Success For User |-----] ^

Objective: Answer whether the login infrastructure, location, or device is common for the user.

[--- IP Space Indicators ---] ^

Source IP: **196.247.229.187**

- First Time Seen for User**

IP Organization: **Packet Exchange Limited**

- First Time Seen for User**

ASN: **58065**

- First Time Seen for User**

Source IP Data Table - Event success 0/163 or 0.00% | Days succeeded 0/11 v

IP Organization Data Table - Event success 0/163 or 0.00% | Days succeeded 0/11 v

ASN Data Table - Event success 0/163 or 0.00% | Days succeeded 0/11 v

[--- Geo-location Indicators ---] ^

Region: **Tallinn, Harjumaa, 10111 EE**

- First Time Seen for User**

Country: **EE**

- First Time Seen for User**

State: **Harjumaa**

- First Time Seen for User**

City: **Tallinn**

- First Time Seen for User**



How to protect your organization from malicious login sources

- Deploy phishing-resistant MFA.
- If FIDO-only factors for MFA are unrealistic, opt for push notifications.
 - If feasible for your environment, configure identity provider policies to restrict access to managed devices.

Pay close attention to detections related to suspicious logins.

- Pay close attention to detections related to suspicious logins. Most technologies will monitor for user behavior that differs from the baseline but these alerts do no good unless they're reviewed and analyzed by security or IT teams.
 - Use vendor technologies that identify or block authentication attempts with previously unseen authentication characteristics, such as impossible

travel, unusual locations for the environment, or a new device for the account.

- Use policies to block access from suspicious logins based on IP address(es), autonomous system numbers (ASN), IP types, or geolocation, and configure pre-auth policies.
- Of note: when analyzing these alerts, remember that malicious logins don't only occur from suspicious locations or VPNs—a lot also come from hosting providers and proxies.
- Use password managers to create, store, and share passwords. These allow users to create unique, complex passwords and store them in a way that keeps them safe from attackers and infostealer malware.

QR code attacks

We found that Microsoft's threat intelligence can reliably identify phishing pages created with PhaaS platforms. Customers using Microsoft's ATP product regularly received alerts for users accessing credential harvesters. However, Microsoft's effectiveness in identifying these pages may have inadvertently contributed to an increase in threat actors' use of QR codes.

Qishing attacks take the activity off endpoints—which Microsoft defenses have visibility into—and put it on mobile devices, which may not be managed with tight security controls. This allows the attacker to bypass endpoint controls.



How to protect your organization from QR code attacks

Microsoft only recently [released new features to address](#) qishing, so we have yet to see the full impact of its response. We believe that to make a real impact, Microsoft needs to block QR codes by default. If Microsoft makes the settings configurable, we recommend organizations opt for the strictest QR code settings.

For organizations that aren't using Microsoft products, it's still important to understand these tactics and ensure users are wary of QR codes just as they might be wary of a URL. And, if an organization uses QR codes for any reason, it should make end-users (customers, partners, employees, vendors, etc.) aware of when and how it uses them.

Targeted attacks

“[The Com](#),” a hacking group including the actors tracked as “Scattered Spider,” was responsible for the largest number of targeted identity attacks this year. These attackers primarily targeted Okta and Microsoft accounts, attempting to abuse password reset mechanisms and MFA push fatigue to gain access.

We observed two main attack tactics: (1) calling into an organization’s helpdesk, and (2) abusing self-service password reset mechanisms.

When calling into the helpdesk the actors impersonated staff and **request** that their passwords be reset. If either the helpdesk requested or the self-service password request attempts were successful, the threat actor sent MFA pushes to the real user. If the user accepted the MFA push, the attacker gained access to the account.

Attacks from these actors are high-volume, too. In one environment, an attacker submitted password reset attempts for more than 100 accounts. In multiple attacks, organizations blocked the attempts thanks to controls that required authentications from a compliant device. In the second half of the year, the threat actor primarily targeted customers in the financial services industry. We expect this style of attack to continue in 2024.

According to Microsoft, these actors [perform thorough research on possible victims](#) and mimic their speech styles. Others in the industry have sounded the alarm on how [AI can make this type of impersonation even easier](#).



How to protect your organization from targeted attacks

Test password reset security controls in your organization. What processes do you have in place for employees calling for a reset? Are those processes required for all employees? Are they followed all the time? Having security controls in place is vital, and regular testing ensures the controls work as expected.

Having security controls in place is vital, and regular testing ensures the controls work as expected.

For often-targeted organizations, we recommend implementing more compensating controls and considering these incidents when evaluating organizational risk. The type and frequency of incidents can provide a powerful argument in requesting funding for additional security controls—especially when well-documented with regular benchmarking and metrics. Specifically, this information can feed directly into risk assessments, which illustrate the potential impact to the organization.

Adversary-in-the-middle (AiTM)

Last year, we noted that threat actors had begun transitioning to AiTM techniques, and we watched use of this tactic grow considerably throughout 2023. AiTM allows threat actors to bypass some forms of MFA, acquire user credentials or session cookies, and gain access by intercepting authentication traffic.

Here's how a typical AiTM attack plays out:

1. **An employee receives a phishing email** with a malicious link or attachment leading to an attacker-controlled website. The website, which acts as a proxy for the genuine website.
2. **Employees submit their credentials** and, after prompting, **provide an MFA challenge answer**.
3. **The attacker proxy relays the requests and responses back and forth** between the employee and the genuine site.
4. The attacker-controlled proxy seamlessly **intercepts the session cookie and redirects the employee** to the application the cookie was set up with.
5. Attackers **import the session token cookies into their web browser and authenticate**, bypassing MFA. Now, the attacker can use the victim's session.





How to protect your organization from AiTM attacks

After a successful attack, organizations must not only reset passwords but also terminate sessions for that user; while the session is still active, the attacker can continue to access the account. In most situations, attackers will attempt to register an MFA device while they have access to the account. We recommend investigating to identify newly registered MFA devices after an account compromise.

Some technology, such as Fast Identity Online 2 (FIDO2) and certificate-based authentication, can stop AiTM attacks. But most organizations haven't quite made the jump to FIDO-only MFA factors—many still use time-based one-time passwords (TOTPs) and push notifications for MFA. While these features raise the bar considerably, gaps remain and AiTM attacks will continue to exploit them. If FIDO2 isn't feasible, consider using other access controls, such as requiring authentication from approved devices.

Some technology, such as Fast Identity Online 2 (FIDO2) and certificate-based authentication, can stop AiTM attacks.

PhaaS platforms have also fed the growth of AiTM, so we don't see these attacks slowing down in 2024. We recommend ensuring your organization automatically terminates sessions and reviews activity performed by an identity-compromised account.

What can you do to detect—and hopefully prevent—these costly attacks? Here's what we recommend for security teams:

- Alert on a user logging in from two different locations. Due to proxying, successful AiTM results in multiple login events—one by the user and one by the attacker. Alert for multiple sessions from the same user with multiple non-mobile operating systems. This is an indicator that something suspicious is under way.
- Alert for potential brute-force push requests.
- Review events marked by end-users as suspicious.

We recommend creating detections for as many stages of a possible incident as is feasible to decrease the possibility of an attack going unnoticed. Examples include:

- Alert on newly added MFA devices.
 - Improve the rate of false positives by tuning out expected network zones.
- [Alert for new Outlook inbox rules created with suspicious names](#). Names that are two to three characters in length, or that contain repeating characters, are generally suspicious.
 - Also watch out for certain keywords, like “payroll,” “malware,” or “virus.”
- Advise employees to be aware of their information in the payroll system and report any abnormal or suspicious activity to the security team. Investigate any unexplained variance in paychecks. Any malicious changes are evidence of a BAC incident.
 - Some human capital management systems allow administrators to require approval for users changing sensitive information (such as direct deposit details). Consider implementing this type of control and other system-specific controls depending on your organization's vulnerability to this type of activity.

Advise employees to be aware of their information in the payroll system and report any abnormal or suspicious activity to the security team.



Cloud platform threats

Credential theft and AWS Cognito attacks

Cloud attacks in 2023: ever upward

A cloud platform threat refers to any incident that occurs in cloud infrastructure, such as resources hosted by AWS, GCP, Microsoft Azure, Docker instances, and Kubernetes clusters (yup, we monitor all of those in depth). We define activity as an incident if an attacker gains at least control plane or data plane access to the environment.

We noted a 72% increase in cloud infrastructure incidents, roughly consistent with what we saw in 2022 and continuing the upward trend we've seen since we began supporting these platforms. Breaking down the numbers further, 96% of the incidents we detected and responded to occurred in AWS, and the remaining 4% were split evenly between GCP and Azure.

96% of the incidents we detected and responded to occurred in AWS, and the remaining 4% were split evenly between GCP and Azure.

This is interesting, since about half of our cloud customers use AWS, about 33% use Azure, and around 17% use GCP. This heavy skew towards AWS aligns with last year's findings, and is likely the result of more AWS security research and auditing tools available for attackers to abuse.

Cloud platform incidents by industry vertical


Chart 6 on page 16 shows that threat actors targeted their cloud attacks on companies in the technology vertical the most in 2023 (nearly 50% of cloud incidents), followed by financial services companies and conglomerates (which we define as a multi-industry company, usually comprising multiple business entities but tracked as a single organization).



WHAT WE SAW IN OUR PREVIOUS ANNUAL REPORT

- Incidents our SOC identified in AWS, GCP, and Microsoft Azure increased by 70% over 2021, underscoring the importance of an effective cloud threat detection and response strategy.
- Of the incidents we spotted in these clouds:
 - 92% occurred in AWS
 - 4% occurred in Azure
 - 4% occurred in GCP
- For context, we monitored roughly four times as many AWS environments as Azure and ten times as many as GCP.
 - **Two major patterns and trends emerged:**
 - **Attackers targeted long-term access keys and service account credentials as a means for initial access.**
 - **Attackers abused public-facing Amazon Elastic Compute Cloud (EC2) instances to perform server-side request forgery (SSRF) or domain name system (DNS) rebind attacks in AWS, or to deploy tooling and malware.**
- Red team exercises or penetration tests accounted for half of cloud infrastructure incidents we spotted in 2022.

Chart 7: Cloud security incidents by vertical industry



These percentages reflect the amount of cloud infrastructure these industries use—that is, the biggest cloud platform users are in the technology industry, followed by financial services. In the following section, we’ll discuss what caused incidents and provide suggestions for organizations securing their own cloud environments.

What we saw

As [Chart 7](#) on page 17 illustrates, exposed credentials were the leading root cause of cloud infrastructure incidents that our SOC observed (42%). Stolen or leaked credentials (or secrets) provide attackers with access to the cloud control plane either through a framework or command-line utility. Publicly exposed or stolen credentials allow attackers to maintain persistent access to the cloud environment with the permissions tied to that identity or role.

Stolen or leaked credentials (or secrets) provide attackers with access to the cloud control plane either through a framework or command-line utility.

Secrets can be exposed through accidental upload to repositories, vulnerability exploitation, or information stealing malware. The secrets users accidentally uploaded to digital repositories were the ones most exposed.


SSRF, which tricks a public-facing web application into exposing sensitive information, was the next most common incident type we saw in 2023 at 28%. Specifically, it tricks AWS EC2 instances into exposing secrets. The main SSRF technique involved attempts to request information about the EC2 instances credentials from `http://169.254.169.254/latest/meta-data/iam/security-credentials/`. (Most cloud instances store their own metadata at `169.254.169.254`.)

[Due to weaknesses in the first version of AWS’s Instance Metadata Service \(IMDSv1\)](#), this SSRF can target AWS instances. Attack mitigation requires organizations to upgrade all SDK and CLI to the latest version to take advantage of IMDSv2.

The third most frequent incident type resulted from use of default credentials (19%), most often abused by adversaries scanning the Internet to deploy crypto miners.

In our [AWS](#), [GCP](#), [Azure](#), and [Kubernetes](#) mind maps, we call out APIs that attackers use or could use with unauthorized access to cloud environments. These investigative resources help us build our detection methodologies and can help your organization in thinking about attacker behavior.

Chart 8: Leading root causes of cloud infrastructure incidents in 2023



Amazon Cognito

Amazon Cognito (also known as AWS Cognito) lets security teams add user sign-up and authentication to applications, but administrators frequently misconfigure the service. In these cases, an attacker can do things like create new accounts with excessive permissions or even access AWS credentials associated with the Cognito Identity Pool.

This year, security researchers released a few [Amazon Cognito auditing tools](#). While these tools help operators identify misconfiguration, threat actors can use these tools as well: we’re seeing more of these Amazon Cognito-focused attacks seemingly enabled by the auditing tools. We expect to see more of this misconfiguration abuse in the coming year because of the high potential impact of misconfigurations.

In the meantime, operators should take the opportunity to secure their environments. [Our recent blog](#) discusses possible misconfiguration in more depth and offers suggestions on monitoring.

Amazon Cognito (also known as AWS Cognito) lets security teams add user sign-up and authentication to applications, but administrators frequently misconfigure the service.





How to protect your organization from cloud threats

Cloud compromise detection opportunities

- For AWS, create a detection methodology around long-term access keys (beginning with AKIA) and short-term access keys (beginning with ASIA) to detect the early stages of unauthorized access to your environment.
- Alert when you see access keys used from unauthorized regions or ASNs that are VPNs, hosting providers, or data centers.
 - Alert on new access keys and key pairs created from anomalous regions or abnormal IP addresses.
 - Monitor for IAM abuse, such as attached policies or new user creation. You can use IP reputation and User-Agent strings to look for higher-risk IAM events.
- Trigger an alert for a high number of failed or unique API calls. This could represent the early stages of discovery that an attacker performs after first gaining access to an environment.
- Review console logins for geo-impossible travel, abnormal source IP address reputation/infrastructure types, abnormal console login times (e.g., off hours), and abnormal user agents.

Other recommendations for protecting against cloud threats

- Follow strong identity management practices.
 - Regularly remove unnecessary keys and rotate access keys.
 - Ensure least privilege for accounts by using role-based security policies.
 - Enforce MFA for access to cloud consoles.
 - Avoid use of static credentials where possible in favor of short-lived, just-in-time access credentials (such as IAM Roles in AWS or service account impersonation in GCP).

- Secure Internet-facing assets.
 - Sometimes people forget to change passwords, no matter how often they're told. To compensate, maintain an inventory of internet-facing assets and ensure the availability of web-access logging. This data aids in investigating and identifying the root cause of an incident, and these logs can help an organization distinguish between a vulnerability exploitation or abuse of credentials.

Sometimes people forget to change passwords, no matter how often they're told.

- Attackers can and will abuse leaked credentials/secrets. To find exposed secrets, configure secret scanning to prevent attackers from abusing sensitive keys. This process identifies exposed secrets before attackers find the vulnerable keys and can also prevent secret exposure in the first place.
 - Services like Github allow organizations to [enable secret scanning](#) (free for public repositories, at a cost for private repositories).
 - Organizations can add custom scans to improve their security postures even more. We use tools such as [Trufflehog](#), which can identify hundreds of secret types from code repositories, S3 and GCS buckets, Docker images, CI/CD platforms, and more, to schedule frequent scans against critical assets.
- If your organization uses Amazon Cognito, perform regular internal audits and include the service in your third-party audits.



Computer-based threats

Malware, vulnerabilities, and ransomware

Before we get into the data, it's important to provide some background. As a managed detection and response (MDR) provider, Expel aims to stop attacks as early in their lifecycle as possible. We then make recommendations to our customer base to help improve their security postures. This need for a quick response and our customers' resilience prevents us from seeing the entire attack chain in most situations. With that in mind, the following analysis of trends primarily consists of interesting, stealthy malware and attacks that bypassed other security controls and required human intervention.

Malware

We separate the malware we identify into categories based on its functionality. This helps us understand the risks introduced into the environment by the malware, the malware's objective, and how to defend against it. [Chart 8](#) on page 20 shows the breakdown of the malware types our SOC observed.

High-risk malware

In over half the malware incidents, the malware deployed presented an immediate and significant risk to the environment—including risks of pre-ransomware activity or exfiltration.


We consider both initial access malware and RATs to be high-risk. Initial access malware (also called "loaders" or "droppers") are lightweight and intended to circumvent defenses before downloading or loading other tools or malware. RATs are aptly named and enable remote access—they can include abuse of legitimate commercial tools, custom attacker-built tools, or a mix of purpose-built tools with nasty and powerful features sold to attackers by an individual or group.



WHAT WE SAW IN OUR PREVIOUS ANNUAL REPORT

- The proportion of organizations attacked by targeted ransomware threat groups in 2022 doubled from 2021.
- 11% of incidents could have resulted in deployment of ransomware had we not been there to intervene—a seven percentage-point increase over 2021.
- When our SOC responded to a deployment or installation of malware, there was a 32% probability the activity was related to ransomware operations.
- Neither industry nor annual revenue provide predictable measures of potential ransomware targeting. This is an everywhere-and-everyone problem.
- The top attack vectors used by ransomware groups to gain initial entry last year were:
 - "Self-installation" techniques on Windows-based computers (97.5% of all ransomware related incidents—an increase of more than 14 percentage points compared to 2021).
 - Exploitation of a software vulnerability on the perimeter (2.5% of all ransomware related incidents—down one-and-a-half percentage points compared to 2021).

Chart 9: Malware observed in 2023 by type



Both initial access tools and RATs provide an attacker access to an endpoint in the environment, setting the attacker loose within the network and creating a high-risk situation. We attribute these attacks to initial access brokers (IABs), who sell access to ransomware gangs or other enterprising threat actors. For this reason, we classify these malware as “pre-ransomware.”

These high-risk threats accounted for 57% of the malware incidents our SOC investigated. Why is this significant? Malware requires a quick response. **Treat actors continue... continue to reduce the time they're in a network before they take action,** with some reports observing ransomware deployed within 24 hours of the attacker's initial access.

These high-risk threats accounted for 57% of the malware incidents our SOC investigated.

The most frequent malware cases we classified as pre-ransomware included:

- Gootloader—17% of high-risk malware incidents
- Qakbot—12%
- SocGhosh—11%

It's worth noting these were also the top pre-ransomware threats we reported in both 2021 and 2022. The skilled actors behind these threats have been active for a while; they'll likely continue to be active for the foreseeable future. An example: the U.S. government [took down Qakbot](#) in August. However, the actor has continued using other malware.





Chart 10: Frequency of high-risk malware by industry



High-risk malware incidents by industry

Among the industries we serve, financial services (19%) and hospitality (17%) experienced the greatest volume of high-risk malware. Based on our review, though, we don't believe attackers specifically targeted these industries. Individual incidents revealed the malware delivery method was more opportunistic—that is, the majority of cases were initiated by drive-by-downloads from infected websites.

Chart 9 sheds light on how frequently attackers target other vertical industries.

Among the industries we serve, financial services (19%) and hospitality (17%) experienced the greatest volume of high-risk malware.



How to protect your organization from high-risk malware

While that all seems dire, there's good news: there are multiple ways to protect your environment from pre-ransomware.

- Take time to understand the attack vectors and how each of them work. Prioritize the computer policies and technology changes that make the most sense for your organization to get the most defense for your buck.
- Simulate activity related to these threat actors through red teaming or threat actor simulation tools (if you're considering a red team exercise, see our blog on [six things to consider before you start](#)). These help ensure your existing technology can detect these threats and help you understand what alerting will look like when the attack happens.
- Leverage threat intel to keep current on what threat actors are up to. They're an innovative bunch and they evolve tactics over time. Members of the security community work hard to expose adversaries' tactics, techniques, and procedures (TTPs), making it easier to reassess your detection coverage regularly

Latent-risk malware

The other category of malware-based incidents we see are what we term latent risks. These malware types don't present an immediate high risk but still threaten an organization's security.

- Infostealers access sensitive data, like passwords, cryptocurrency wallets, and other information stored in user's browsers or on their devices.
- Banking Trojan steals or intercepts financial information from the user's browser.
- Browser hijacker redirects the user to advertisements or malicious websites.
- Coinminers use the resources of a computer or server to generate or mine for cryptocurrency on behalf of the attacker.

Infostealers represent the most serious of these latent risks. Like high-risk malware threats, they provide a wide range of options for an attacker. At a minimum, infostealers exfiltrate passwords to the attacker, who can use them for fraudulent purposes or sell credentials belonging to VPNs to other attackers looking for an organization to target (or who are specifically targeting your organization already).

While infostealing malware can be targeted, threat actors most frequently deploy it opportunistically.

While infostealing malware can be targeted, threat actors most frequently deploy it opportunistically. Adversaries often take advantage of a user looking to download software. "Cracked" software, or software that generally requires a license but attackers have modified to circumnavigate that requirement, represents the highest risk. In these situations, users are more likely to act against their better judgement in downloading software from a sketchy website to avoid buying a paid version.

We also saw a lot of infostealers downloaded from malicious ads (also known as malvertising), which appear at the top of search results and often imitate productivity or IT software. Clicking the ad directs the user to a cloned version of the legitimate software's website. When the software downloads and runs, the malware executes the malicious payload—sometimes even installing the legitimate software to avoid suspicion. This behavior also heightens the risks surrounding shadow IT. Speaking of which...

A risk of shadow IT

Shadow IT refers to the use of hardware or software that an organization's IT or security team hasn't vetted. The appetite for risk determines the degree of vetting required.

Throughout the year, we frequently saw attackers leveraging both search engine ads and SEO poisoning to guide users to download malicious payloads. The main types of malware used in these attacks are RATs and infostealers. Common "lures" include software used by IT teams (such as Advanced IP Scanner, Nmap, and WinSCP) and productivity software (such as Notion, Notepad++, and PDF manipulation tools). These attacks target both Windows and MacOS systems.

Throughout the year, we frequently saw attackers leveraging both search engine ads and SEO poisoning to guide users to download malicious payloads.

Shadow IT attacks became a fad for threat actors in 2023. Google and Microsoft actively work to combat the tactic, but we expect malvertising to remain popular into 2024 (and ultimately, for as long as it works).

Both tactics—malvertising and SEO poisoning—often use [anti-analysis techniques](#). Attackers also regularly use different techniques to implement anti-analysis to prevent automated analysis of downloaded files, such as inflating the size of the malware.

Inflated infostealers

Inflating malware size—a tactic that involves an attacker adding junk data to an executable file to make it very large—was popular in 2023. For example, an attacker might take a file that's 300 KB in size and add random bytes to increase the size dramatically, making the file 900 MB. These extra bytes prevent analysis from antivirus tools, EDR tools, or other automated analysis tools, like sandboxes.

Most public and private sandboxes have a size limit, typically between 100 and 400 MB. These limitations make it difficult for security tech to detect suspicious file attributes. Organizations without a malware analyst may be completely unable to triage the file. While this technique isn't new, more attackers use it and build it into their platforms to make it easy for other bad actors to take advantage.



How to protect your organization from latent-risk malware

- Policies that require effective vetting before users download anything are important, but so is a healthy organizational culture and awareness. **Do you have** a process in place for users to place requests for helpful tools? Do users feel like their software requests will be filled quickly? **What if users are part of the IT department itself—do they need to follow the same process?** A healthy security culture helps users make choices consistent with the software policy
- Review logs and a user's browser history to understand what really happened. Anti-analysis often thwarts security or IT teams' attempts to recreate the behavior.
- Attackers also **regularly** use different techniques to **implement anti-analysis to** prevent automated analysis of downloaded files, such as inflating the size of the malware. **Security analysts** can circumvent these inflation techniques using tools like **debloat**, which which remove the added junk bytes, allowing for manual or automated analysis and triage to understand the malware's behavior more easily.



Vulnerabilities

The top vulnerabilities our SOC identified based on incident frequency included:

1. Progress' MOVEit Transfer (CVE-2023-34362)
2. Adobe ColdFusion (CVE-2023-26360)
3. Citrix NetScaler ADC, also known as CitrixBleed (CVE-2023-4966)

All of these vulnerabilities were zero-days—that is, adversaries exploited them before vendors were aware and could fix the problem. Zero-days are terrifying because by the time the vendor finds them, the attacker already fully understands the vulnerability and is exploiting it—putting the vendor and security community at a significant time disadvantage as they scramble to catch up. Further, it

takes time **for the community** to determine how long the vulnerabilities have existed and been exploited.

However, security teams can identify zero-day exploitation when it happens thanks to defense-in-depth and by understanding common attacker tactics. **As operators, we** continue to learn lessons from these incidents, allowing us to better proactively address risk and enhance resilience every time a zero-day occurs.

Zero-days emphasize the importance of defense-in-depth and layers of controls. Properly configured and implemented security controls—or lack thereof—directly affect the impact of these vulnerabilities on an organization. Tools such as EDR and web application firewalls (WAFs) let operators observe activity associated with exploitation and quickly act. Without these, software and hardware become a single point of failure with devastating consequences.



How to protect your organization from vulnerabilities

To reduce risk to your external-facing assets, consider:

- Constantly scrutinize all external-facing assets with frequent (as operationally feasible) scans that identify vulnerabilities and configuration risks. Additionally, only enable business-essential ports and services.
- Security and IT teams should collaboratively monitor external-facing infrastructure for gaps in detection and defense. Use metrics to show work completed remediating incidents and bolstering defenses: these metrics can prove valuable at the executive level to show work the security team is performing and support requests for more resources as needed. This data is valuable for risk assessments and tracking risk mitigation.
- Attackers can't easily leverage vulnerabilities, but if they're severe enough, then attackers tend to move quickly toward exploiting exposed assets. If the sheer number of vulnerabilities overwhelms your team, the [Expel Vulnerability Prioritization](#) solution can help sift through alerts, automatically surfacing the highest-risk vulnerabilities to help teams focus on the threats that matter most.

Ransomware

Threat actors who target businesses with ransomware may steal data, threaten to leak data, or further harass the targeted organization.

Ransomware happens at the end of an attack lifecycle. Deployment depends on other threats we talk about in this report: identity compromise, RATs, credential theft via infostealers, exploitable vulnerabilities, phishing, and more. When organizations don't monitor and mitigate these other threats, the risk of an attacker gaining access to systems and ransoming critical data and computer systems grows.

In the next section, we discuss trends and lessons learned over the year about monitoring and mitigating the threat of ransomware.

Ransomware gangs

Before we get into tactics, it's important to mention the gangs associated with ransomware. Active gangs change from year to year, and over time they may dissolve, change names, disappear, or be taken down by law enforcement (only to then re-emerge in a different form later). Gangs have different organizational structures, and they differ in their preferred tactics. Therefore, the following section talks broadly about observed tactics and information gathered from threat intelligence.

For information on specific ransomware gangs, we recommend following the [Cybersecurity & Infrastructure Security Agency's \(CISA's\) news publications](#).

Ransomware actors often buy stolen credentials, which they use to leverage the organization's own VPN or a remote access tool to access the environment.

Initial access

Throughout the year, attackers used the following primary tactics to access victim's environments:

- 1. Leveraging initial access malware.** As discussed previously, initial access malware affords attackers a door into an environment. Access to that door is managed by initial access brokers (IABs); ransomware gangs commonly buy from IABs and often invest in the development and deployment of initial access malware.

- 2. Exploitation of software vulnerabilities on perimeter assets.** Attackers may leverage a vulnerability to let themselves into the environment. Some ransomware gangs prefer this tactic and use it almost exclusively.
- 3. Use of compromised credentials to access the environment.** As we've discussed, attackers find many ways to compromise credentials. Ransomware actors often buy stolen credentials, which they use to leverage the organization's own VPN or a remote access tool to access the environment.

Security tech provides opportunities to detect an attack at multiple stages. However, when visibility is limited, there's a high risk of compromise. These low-visibility situations often allow an attacker to act without being seen or mitigated. It's essential to ensure you have sufficient monitoring in place for all assets in your environment.

Pre-ransomware initial access

IABs, which trick users into running malware, were the top access method for ransomware actors among our customers in 2023. To defend against this tactic, organizations should look for mitigation techniques for the malicious use of the following file types:

1. JavaScript
2. Executables
3. LNK


In 2023, bad actors trended toward script-based files for pre-ransomware malware. JavaScript made up the highest volume of files at 39%, but we also saw lots of other scripting types, including executable files (EXE) (20%), LNK (12%), VBS (7%), and others, as illustrated by [Chart 10](#) on page 25.

JS and VBS are text files that receive less scrutiny from automated analysis mechanisms than EXE files. Unlike EXE, [attackers can hide their functionality](#) among benign content or by command obfuscation. The malicious behavior is only observable when the file is executed by another application.

By default the native Windows program, wscript.exe, executes these scripts when a user double-clicks the files. Attackers use this setting to their advantage.

The power of secure by default: At the very end of 2022, we saw the beginning of [a new trend toward leveraging OneNote files](#). In January 2023, Microsoft pushed a patch to slow down the exploitation, and ultimately implemented much stronger controls in March, virtually killing the technique. This not only illustrates how Microsoft can eliminate a threat

Chart 11: Top file types used to deploy pre-ransomware malware in 2023



vector, but it also shows that if organizations adopt tight controls, they also have the power to fully stop a threat vector.

Vulnerability exploitation

While some ransomware gangs rely on initial access tools, others access environments by targeting vulnerabilities in external-facing appliances. The ransomware gang Play, which uses this technique, is [known to exploit vulnerabilities in the operating system of Fortigate’s firewalls \(known as FortiOS\) and Microsoft Exchange Server](#).

Understanding how ransomware groups leverage vulnerabilities highlights the importance of knowing what vulnerabilities exist in the environment. Defense against ransomware, then, also involves defense against vulnerability exploitation.

Access by compromised credentials

Throughout 2023, multiple ransomware groups leveraged compromised credentials to access customer environments. Infostealers, vulnerabilities, social engineering, and other methods can all lead to compromised credentials.

One ransomware gang in particular, Akira, [targeted VPNs which didn’t have MFA enabled](#). This insecure configuration let the threat actor try multiple methods (like brute force, password spraying, or buying stolen credentials) to access accounts. Once attackers identified valid credentials, they gained initial access and then worked to compromise the network.

Earlier in this report, we discussed The Com. These actors used social engineering to acquire valid credentials to then access the victim’s environment. Once they had access, the group deployed BlackCat ransomware against the victim’s networks.

These three initial access methods highlight where an organization needs to focus its defenses. Defending against encryption is too late in an attack to protect an organization, but defending against initial access methods can help deflect the overall threat.





How to protect your organization from ransomware

Top initial access patterns

Some behaviors are much more suspicious than others. Alerting on behavioral patterns like the following helps spot all kinds of malicious activity, not just ransomware gangs trying to infiltrate your network.

Endpoint behavioral patterns

Alert when...

- A scripting process other than PowerShell (like wscript.exe) launches a PowerShell process with encoded commands.
- The Microsoft HTML application host utility, mshta.exe, loads a command line and makes network connections.
- Scripting processes like wscript.exe or cscript.exe do the following:
 - Execute a .vbs, .vbscript, or .js file.
 - Initiate an external network connection.
 - Spawn a cmd.exe or PowerShell process.
- PowerShell executes a base64 encoded command and the process initiates an external network connection.

Network behavioral patterns

Alert when...

- PowerShell makes a network connection to a rare external destination.
- A high volume of Kerberos ticket requests and DNS requests for internal host names come from the same source.

When building behavioral detections, start with a threat hunting-based approach to answer questions like, “How often over a 30-day period do we see this type of behavior?” or, “When we see PowerShell initiate an outbound external connection, what are the typical destinations?” Then apply your insights and enable new network-based behavioral alerts for activity that’s not normal for your organization.

Other recommendations for protecting against ransomware:

- Configure JavaScript (.JS, .JSE), Windows Script Files (.WSF, .WSH), and HTML for application (.HTA) files to open with Notepad. By associating these file extensions with Notepad, you mitigate a common entry point for malware.
 - If you can’t use Notepad, change the wscript.exe default setting so it doesn’t auto-execute scripts when a user double-clicks the files. You can do this using a group policy object (GPO).
- Unregister ISO file extensions in Microsoft Windows Explorer. Once this is done, Windows will no longer recognize ISO files and double-clicking won’t result in program execution.
- Don’t expose remote desktop protocol (or any other service you don’t need to) directly to the Internet.
- Keep an inventory of exposed assets, and update the list and assets regularly.
- Prioritize vulnerabilities to patch in your environment. Have a flexible program that allows you to patch out of band when critical vulnerabilities are identified.

One last piece of advice: **have a ransomware incident response (IR) plan and test it.** A real-life attack isn’t the best time to test your plan. A great one emphasizes roles and responsibilities, communication, reporting, how to handle data, and how to prepare for the emotions your team will experience. Stress test your plan regularly—we recommend once a quarter—to make sure everyone knows what to do when a bad thing happens. (If you need a starting point—or just want a fun and familiar structure for your IR tabletops—our updated [Oh Noes! role-playing game](#) can help.)

One last piece of advice: have a ransomware incident response (IR) plan and test it.



Phishing threats

Phishing and credential harvesters

You may have noticed many of the risks mentioned in this threat report use phishing as an attack vector. That's because it's a means of targeting user identity and introducing multiple types of malware into a victim's environment. Similar to how web servers must be exposed to the Internet, it's commonly necessary for employees to receive and open emails to perform their jobs—and there's no way around it. Both web servers and emails pose inherent risks for organizations. But, like all risk, this must be accepted or mitigated.

One option we provide for mitigation is to triage emails submitted by end users in customer organizations. We're then able to investigate whether other users in the environment received the same email or if there are signs anywhere in the environment that the malware was successful.

Here's how it works: end users submit emails they deem suspicious and the Expel Phishing Team triages and investigates to determine if they're malicious, unwanted spam, or legitimate communications. Our team can action malicious emails on behalf of our customers, and this triage not only relieves the burden from an organization, but affords us unique visibility into the various phishing attacks launched against a range of organizations.

The emails triaged by our team made it past other security tools, such as email gateways, and landed in a user's inbox.

Phishing incidents by industry

Among the industries we serve, hospitality saw the highest volume of targeted phishing attacks at 55%. Hospitality is followed—but not closely—by travel (12%), technology (9%), financial services and healthcare (tied at 5%). See [Chart 11](#) on the next page for a full breakdown.




WHAT WE SAW IN OUR PREVIOUS ANNUAL REPORT

- **16% of phishing emails submitted by customer employees proved malicious. Of those, 88% were credential harvesters. Malware, banking fraud, and gift cards scams combined to account for the other 12% of malicious emails.**
- **The PDF file format accounted for 84% of the malicious attachments spotted by our phishing team. The remainder compromised other historically common file types (.ZIP, .DOCX, .XLSX). Attackers don't weaponize the majority of these files with an exploit. Instead, they embed an evil link in a PDF, nest an executable in a ZIP file, or create a macro that must be "run."**
- **Organizations that frequently correspond with a high number of third-party vendors were most likely to experience a successful phishing attack.**



Chart 12: Phishing attacks by industry in 2023



If you recall, we noted earlier in the report that hospitality, technology, and financial services also made it into the list of the top industries where we identified the most high-risk malware and identity incidents.

Top phishing trends

Over the year, attackers most frequently used HTML/HTM attachments to direct targeted victims to a credential harvester (see image below for an example). In their simplest form, these documents only contain enough code to load the harvesting website. The sites themselves frequently have bot protection from CloudFlare or other anti-analysis settings to prevent review of the malicious domain.

Qishing

Throughout 2023, our analysts noted a rise in the abuse of QR codes for phishing (aka "qishing"). Similar to other tactics, such as AiTM, the growth of this attack was made possible by

popular phishing platforms, such as DadSec/Phoenix, which added QR code support for its clients.

While easy to use, QR codes increase the danger associated with phishing. With a URL, a user can visit the malicious domain using the organization's endpoint, giving operators the opportunity to block connections using multiple technologies. However, with a QR code, the activity moves off of the workstation and onto the user's mobile device.

These attacks are, again, best mitigated by security-in-depth, ensuring security controls exist in different stages. Not all users will accurately identify phishing, so security or IT teams should review suspicious login activity for users. Malicious logins may go undetected, so security teams need to create and review suspicious activity associated with accounts, such as suspicious inbox rules. Security-in-depth allows for multiple opportunities to detect malicious activity.

Image 3: HTML contents of a file that loads a credential harvester

```
1 <!DOCTYPE html>
2 <html lang="en">
3 <head>
4   <title>Bootstrap Example</title>
5   <meta charset="utf-8">
6   <meta name="viewport" content="width=device-width, initial-scale=1">
7   <link rel="stylesheet" href="https://maxcdn.bootstrapcdn.com/bootstrap/3.4.1/css/bootstrap.min.css">
8 </head>
9 <body>
10
11 <div class="container">
12 <script>
13 window.location.replace("https://astreahoodgrove-quartzelec.onLine/?e=Y2hhcmxlcY5maW5zdGVyQHQJ1Z3JhdHMuY29t");
14 </script>
15 </div>
16
17 </body>
18 </html>
```



How to protect your organization from phishing

Security control recommendations

- Make sure you're running MFA wherever possible and using phish-resistant FIDO security keys to significantly reduce the risks associated with credential theft.
 - Consider deploying a secure email gateway (SEG) to monitor incoming and outgoing emails for signs of an attack.
 - Use email anti-spoofing controls such as DMARC, SPF, and DKIM.
 - Organizations using Microsoft products should also consider [Microsoft's tool for handling QR codes](#), which is built into Microsoft Defender for Office 365. The exact details of their capability isn't clear yet, but we recommend keeping up to date with developments and using the most secure settings.
- Organizations in heavily targeted industries should take extra security measures, including phishing training, strong endpoint defenses, and regular reviews of the environment based on user-identified malicious emails.
 - Triaging user-submitted emails can also help uncover compromise within an organization.
 - Remind users to treat QR codes with the same suspicion they treat URLs. Never scan a code unless it's received through a trusted, verified source. Inform your employees what forms of communication you will and won't use, and whether they should expect QR codes from leadership or internal teams.
 - Open-source decoding tools, such as [zxing.org](#), also exist to help both end-users and security teams alike—telling you where the QR code leads before you follow the link.

Other recommendations for protecting against phishing threats

- As we head into an election year in the United States, organizations should reinforce the importance of phishing-related training—including how to recognize suspicious or unknown links and spot common lures. Threat actors will exploit users' political leanings and loyalties to generate convincing phishing and spear phishing attacks, likely with the help of AI.
 - Additionally, educate specific business units on the phishing campaigns that might target them. For example, doctors might see medical-themed lures that prey on emerging health concerns, while finance teams might encounter financial-themed campaigns, with "URGENT:INVOICES" subject lines. Recruiters may see résumé-themed phishing lures.

Remind users to treat QR codes with the same suspicion they treat URLs. Never scan a code unless it's received through a trusted, verified source.



Annual spotlight

Infostealer campaign targets hospitality

Since 2022, we've observed an infostealing campaign targeting our hospitality customers with the goal of gaining administrative access to sites like Booking.com to commit fraud.

Campaign overview

The campaign has developed over time, but it generally looks like this: an attacker uses a Gmail account to request information about a booking, ask for help, or to lodge a complaint. There's a strong likelihood employees will interact with these emails—engaging with communications like this is the employee's job, after all. Instead of an attachment, the email contains a link to a file storage service such as Dropbox, Google Drive, and Mega.nz.

In most situations, the files in the storage service are a file archive (ZIP, RAR, etc.), which is password-protected to prevent the storage provider from scanning its contents. When the user opens the file, the archive typically contains an inflated EXE file, which as we outlined above, is an infostealer.

Though this campaign targets hospitality companies, understanding it is extremely valuable to organizations in other industries too, as a common threat vector infostealing malware can impact any organization. It's important to consider the potential implications and impacts for your industry. What credentials, if stolen, could harm your organization and/or your customers?

Though this campaign targets hospitality companies, understanding it is extremely valuable to organizations in other industries, too.



How to protect your organization from infostealers

We recommend using a dedicated password manager. Password managers generate and store passwords in an encrypted format until they're needed; they also facilitate secure password sharing among team members, preventing credential compromise, and thus helping block infostealers.





Looking ahead to 2024...

Thoughts from our Operations teams



Greg Notch

Chief Information Security Officer

Identity has been and will continue to be the frontier for risk. With location and infrastructure control no longer core places where security controls are added, access and identity controls are the new firewall. Adding to this complexity is the rise of LLMs and generative AI technology, making the determination and re-validation of someone's identity much more difficult (for example, onboarding a remote employee or doing a password reset in a world with deepfake video tools). We're just starting to see the class of problems where third parties are given access to company systems with no real way to validate identity.



Daniel Clayton

VP, Security Operations

Bad actors aren't hindered by a concern for the risks associated with AI, so their use of the technology for cybercrime will significantly outpace security team adoption. Adversarial use of AI will supercharge social engineering and a new generation of spear-phishing attacks. The election cycle and emotive geopolitical situations provide a particularly rich breeding ground for disinformation. As such, fast-evolving, intricately personalized, and uniquely convincing AI-driven spear-phishing attacks will mean more compromised credentials and stolen sensitive information than ever before.

“Adversarial use of AI will supercharge social engineering and a new generation of spear-phishing attacks. The election cycle and emotive geopolitical situations provide a particularly rich breeding ground for disinformation.”



Ray Pugh

Director, Security Operations

Tried-and-true methods historically used by attackers will remain in use as long as they're effective. Over the past several years, we've seen clever updates to their approaches (QR code links in phishing emails, for example), and we're likely to see other incremental evolutions this year. It's unlikely that attackers will fully pivot away from email any time soon, but other routine forms of communication (Slack, Zoom, and other collaboration tools) may rise in popularity and add more diversity as vectors for attackers to achieve their objectives.



Ben Brigida

Director, Security Operations

Everyone is looking for ways to maximize profits, and attackers are no different. I believe attackers are going to continue to look for ways to blend attacks to achieve the most profitable outcome. An example might include trying to extort individuals (even after a ransom is paid), and leaving some ransomware or backdoors to return to the victims' systems later. This is already starting to happen in some cases, but it's possible it becomes standard operating procedure for more threat actors in the near term. It's more important than ever to stop the attacker before they complete their mission, because they'll pivot to completing a second and third mission in the same environment.



Steve Edwards

Director, Detection and Response

Phishing will continue to remain prevalent as both an effective and inexpensive means for attackers to compromise organizations. This isn't really a technical vulnerability, but a human vulnerability. Phishing, at its core, is simply one human lying to another. Whether the adversary's desired outcome is credential harvesting or deploying malware, the crux of the attack is convincing targets to do something they know they

shouldn't. Since there is no Patch Tuesday for the human OS, security leaders will need to continue to find ways to allow people to fail safely. Hardened credentials, modern EDRs, and closely monitoring for signs of compromise continue to be the critical "basics" for building a security program.



Christine Billie

Detection and Response Manager

I firmly believe that looking forward through the windshield should be prioritized for any security organization, but I'm always wary of what can happen if we fail to periodically check the rearview mirror. This is especially important given the recent trend of attackers recycling "old" attack vectors that newer analysts may not have seen yet—and are least expecting. For example, our managed phishing service saw a [homoglyph attack](#) this year—and it's unlikely that many of our customers' security teams have seen this tactic before. We also recently saw an attempted HTML injection attack from a threat actor who was hoping that email subject lines would not be sanitized or converted to plain text prior to being ingested and parsed. While both of these are considered "old school" attacks by seasoned security professionals, it's a fun opportunity for SOC managers to take trips down memory lane to teach analysts about attack vectors that aren't necessarily new or trendy, but are still potentially lethal.



Oscar De La Rosa

Detection and Response Analyst

Last year was AI's coming-out party. In 2024, we should expect to see it play a larger role, as it can streamline some of the attacker's infrastructure. From enabling better social engineering attacks (phishing, smishing, vishing) or just helping with the increased deployment of malicious activity, AI is already out there and we can see it being integrated within more attack flows.

"From enabling better social engineering attacks (phishing, smishing, vishing) or just helping with the increased deployment of malicious activity, AI is already out there and we can see it being integrated within more attack flows."

I also think that adaptation for security controls and tools should follow, helping operators with all sorts of critical tasks. Implementation will be paramount for minimizing dislocations for security professionals and end users. With all of that being said, we should keep in mind the basics tenets of proper security practices for end users. Ongoing monitoring of conditional access, MFA, and application-review policies are great things to always have up to date within any system.



Aaron Walton

Threat Intelligence Analyst

Perhaps they're not as flashy as AI, but the reality is that the main threats for 2024 are pretty much the same threats—and the same people—from years past. Year-over-year, we see the same groups and individuals execute attacks successfully. These adversaries have built infrastructure that allows them to carry out attacks, and they have the skills to pull it off.





2023 at Expel

Headlines, accolades, and research

AWARDS

- [Fortune Cyber 60](#)
- [2023 Deloitte Technology Fast 500™](#)
- [CRN® Security 100](#)
- [CRN® Channel Chiefs 2023](#)
- [CRN® Partner Program Guide 2023](#)

INDUSTRY RECOGNITION

- [Expel named a Leader in The Forrester Wave™: Managed Detection And Response, Q2 2023](#)

RESEARCH

- [The SANS Institute - Frameworks, Tools, and Techniques: The Journey to Operational Security Effectiveness and Maturity](#)
- [Enterprise Strategy Group \(ESG\) - Analyzing the Economic Benefits of Expel's Managed Detection and Response Services](#)
- [The Cloud Security Alliance \(CSA\) - Security-Enabled Innovation and Cloud Trends](#)

NEWS

- [Visa Enters Strategic Partnership with Expel to Help Clients Manage Cybersecurity Risk](#)
- [Expel Demonstrates Partner-first Commitment with Revamped Partner Program](#)
- [Expel Appoints Seasoned Hyper-Growth Chief Product Officer to Leadership Team](#)

PRODUCT AND FEATURE LAUNCHES

- [Expel Announces New Vulnerability Prioritization Solution](#)
- [Expel Advances Leadership in Cloud Security with MDR for Kubernetes](#)
- [Expel Workbench History = unparalleled MDR transparency](#)
- [How we built it: the app that gives our analysts more time to fight cyber evil](#)
- [Okta cross-tenant impersonation: a new Expel detection](#)

GREAT PLACE TO WORK®

- [Expel Named to Five 2023 Best Workplaces Lists by Great Place to Work®](#)

EXPEL QUARTERLY THREAT REPORTS 2023

- [Q1 2023](#)
- [Q2 2023](#)
- [Q3 2023](#)

WANT TO LEARN MORE ABOUT EXPEL?

- [Understand how Expel helps secure and grow your business](#)
- [Learn how to simplify your multi-cloud security with Expel](#)
- [See why Expel was named a Leader in The Forrester Wave™: Managed Detection And Response, Q2 2023](#)
- [Subscribe to our blog](#)
- [We're hiring! Find the right role for you](#)
- [See what Expletives say about working at Expel](#)

Reference highlights

Sources consulted and our author

Sources

- FBI. (2023, August 29). FBI, Partners Dismantle Qakbot Infrastructure in Multinational Cyber Takedown. [FBI.gov](https://www.fbi.gov).
- Gandhi, U. (2023, December 12). Protect your organizations against QR code phishing with Defender for Office 365. [Microsoft.com](https://www.microsoft.com).
- IC3. (2022). Federal Bureau of Investigation Internet Crime Report 2022. [ic3.gov](https://www.ic3.gov).
- Lakshmanan, R. (2023, May 13). New Phishing-as-a-Service Platform Lets Cybercriminals Generate Convincing Phishing Pages. [The Hacker News.com](https://www.thehackernews.com).
- MacCarthaigh, C. (2019, December 19). Add defense in depth against open firewalls, reverse proxies, and SSRF vulnerabilities with enhancements to the EC2 Instance Metadata Service. [Amazon.com](https://www.amazon.com).
- Microsoft Incident Response. (2023, October 25). Octo Tempest crosses boundaries to facilitate extortion, encryption, and destruction. [Microsoft.com](https://www.microsoft.com).
- Microsoft Threat Intelligence. (2023, November 30). Microsoft has detected Danabot (Storm-1044) infections leading to hands-on-keyboard activity by ransomware operator Storm-0216 (Twisted Spider, UNC2198), culminating in the deployment of Cactus ransomware. In this campaign, Danabot is distributed via malvertising. [X](https://x.com).
- Microsoft Threat Intelligence. (2023, December 15). Microsoft has identified new Qakbot phishing campaigns following the August 2023 law enforcement disruption operation. The campaign began on December 11, was low in volume, and targeted the hospitality industry. Targets received a PDF from a user masquerading as an IRS employee. [X](https://x.com).
- Proofpoint. (2023, May 12). Crime Finds a Way: The Evolution and Experimentation of the Cybercrime Ecosystem. [Proofpoint.com](https://www.proofpoint.com).

- Santos, O. (2023, October 24.) Akira Ransomware Targeting VPNs without Multi-Factor Authentication. [Cisco.com](https://www.cisco.com).
- Siddiqui, Z., Bing, C., & Satter, R. (2023, November 15). FBI struggled to disrupt dangerous casino hacking gang, cyber responders say. [Reuters](https://www.reuters.com).
- Tidy, J. (2023, December 21). Lapsus\$: GTA 6 hacker handed indefinite hospital order. [BBC](https://www.bbc.com).

Meet the author



Aaron Walton
Analyst, Threat Intel

Aaron Walton is a Threat Intel Analyst at Expel. In this role, he monitors threat actor trends and behaviors to support Expel's operations. **He recommends following @ExpelSecurity on X (formerly Twitter) for articles published by him or his team.**



ABOUT EXPEL

Expel helps companies of all shapes and sizes minimize business risk. Our technology and people work together to make sense of security signals—with your business in mind—to detect, understand, and fix issues fast.

Powered by our security operations platform, Expel offers managed detection and response (MDR), remediation, phishing, vulnerability prioritization, and threat hunting. For more information, visit our [website](#), check out our [blog](#), or follow us on [LinkedIn](#).

expel.com/contact-us