



Great expeltations

Cybersecurity trends and predictions

2023

Contents

A message from merk (aka our CEO) 3

Too long, didn't read: the TL;DR from our SOC 4

Identity threat detection trends 6

Cloud platform security trends 16

Ransomware 23

Phishing 29

Looking ahead to 2023..... 32

About Expel 33

Appendix..... 34

A message from merk (aka our CEO)

Welcome to Great eXpeltations 2023, the second iteration of Expel's annual cybersecurity trends report.

For the past year, we've shared quarterly threat reports (QTRs) to provide data and insights on the attacks we're seeing, how to spot them, and the top ways you can protect and enable your organization. Great eXpeltations is the amalgamation of those reports, plus new findings for the fourth quarter of 2022 (Q4) and predictions for what's to come in 2023.

From the jump, transparency has been a core value here at Expel. Now, as part of our commitment to walking-the-walk, we're sharing insights and lessons learned from our customers to help bring the defender community together against our common adversaries.

This report, co-authored by our security operations center (SOC) leadership team, aggregates data from the incidents we investigated in 2022. It covers our full customer base, from small and midsize companies to enterprise organizations at every phase of the security journey across every industry we serve.

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 expect for 2023. We cap it off with some thoughts and predictions for cybersecurity this year, courtesy of our SOC leadership team.

Whether you've followed along with every quarterly report or this is your first time here, the goal remains the same: to translate the events we detect into security strategy for your organization by sharing how attackers got in, and how we stopped them.

Good hunting everyone.

merk



Dave Merkel
CEO, Expel

“From the jump,
transparency
has been a core
value here at
Expel.”

Too long, didn't read: the TL;DR from our SOC

The trends and predictions we discuss in this report are based on incidents our SOC identified through investigations into alerts, email submissions, and threat hunting leads spanning January 1 to December 31, 2022. We relied on a combination of time series analysis, statistics, customer input, and analyst instinct to identify these key insights.

Our goal: delve into the data to identify trends and patterns in the security events we're seeing and translate them into a security strategy 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.

But before we get into the details (or if you're short on time), here's the bottom line:

Representing 50% of all incidents, business email compromise (BEC) remained the top threat to our customers. This is consistent with 2021.

- Of the BEC attempts we identified: more than 99% were in Microsoft 365 (M365—previously known as Office 365, or O365) and fewer than 1% occurred in Google Workspace.
- Fifty-three percent of all organizations experienced at least one BEC attempt, and one organization was targeted 104 times throughout the year.

Threat actors started moving away from authenticating via legacy protocols to bypass multi-factor authentication (MFA) in M365. Instead, they adopted frameworks such as Evilginx2, facilitating adversary-in-the-middle (AiTM) phishing attacks to steal login credentials and session cookies for initial access and MFA bypass.

- FIDO2 (Fast ID Online 2) and certificate-based authentication stop AiTM attacks. However, many organizations don't use FIDO factors for MFA.

Threat actors targeted Workday to perpetuate payroll fraud.

- Starting in July 2022, our SOC identified BEC attempts, across multiple customer environments, targeting access to human capital management systems—specifically, Workday. The goal of these attacks? Payroll and direct deposit fraud.
- Once a threat actor accesses Workday, they modify a compromised user's payroll settings to add the attacker's direct deposit information, redirecting the victim's paycheck into the attacker's account.

MEET THE AUTHORS



Jon Hencinski
VP, Security Operations

[@jhencinski](#)

Team builder and people leader on a mission to help the industry "SOC" the right way. Believes a SOC *can* be a great place to work—but highly effective management is required!



Ben Brigida
Director, SOC Operations

[@The_Real_BenB](#)

Passionate about helping customers, employees, and coworkers. Focused on finding the bad guys and leading with empathy.



Ray Pugh
Senior Manager,
Security Operations

[@Rpughlll](#)

SOC work is both exciting and rewarding when you're surrounded by the right people and focused on a noble mission. Fueled by a passion for helping others.

Additional contributors: David Blanton, Brian Badorrek, Tucker Moran, Myles Satterfield, Dan Whalen, and Simon Wong

- The lesson? Enforce [MFA within Workday](#) and [implement approval workflows](#) for changes to direct deposit information.

Eleven percent of incidents could have resulted in deployment of ransomware had we not intervened—a seven percentage point increase compared to 2021.

- As Microsoft continues making it easier for organizations to [block macros in files downloaded from the internet](#), ransomware threat groups and their affiliates are abandoning their use of visual basic for application (VBA) macros and Excel 4.0 macros to gain initial entry to Windows-based environments. Instead, ransomware operators opt to use disk image (ISO), short-cut (LNK), and HTML application (HTA).
- Attackers used zipped JavaScript files to gain initial access in 44% of all ransomware incidents.
- ISO files were used to gain initial access in 12% of ransomware incidents. This attack vector didn't make our list in 2021.
- Nine percent of all ransomware incidents started from an infected USB drive.

Attackers targeted long-term access keys and service account credentials to gain initial access to public cloud environments.

- Of the cloud security incidents we responded to, 20% resulted in deployment of cryptocurrency miners and 13% resulted in deployment of publicly available commodity malware.

Six percent of business application compromise (BAC) attempts used push notification fatigue to satisfy MFA.

- Push notification fatigue occurs when attackers send repeated push notifications until the targeted employee “authorizes” or “accepts” the request, allowing the attacker to satisfy MFA.

Credential harvesters accounted for 88% of malicious email submissions. Credential theft via phishing continues to grow, with identity the main focus of today's attacks.

- The top subject lines in malicious emails that resulted in an employee click or compromise were, “Incoming Voice Message,” “Checking in,” and “Voice Mail Call received for <user's email>.”
- Our data shows that social engineering themes that are actionable, time-sensitive, and financially driven are most successful.



Identity threat detection trends

Targeting emails, application data, and payroll settings

What are identity threats?

Complicated jargon runs rampant in the security industry, so let's talk about what we mean when we say **identity threats** in the context of this report: *attackers attempting to gain access to a user's identity to perpetuate fraud*. That is to say, this section will *not* cover identity theft where attackers compromised a virtual private network (VPN) account to steal intellectual property or to deploy ransomware.

Attackers use BEC and what we call BAC as a means to access those identities, so we have historically grouped our data findings separately to illustrate the difference in these entry points. But, BEC and BAC really fall under the same umbrella—identity fraud—which includes things like tax form (aka W2) scams, romance scams, real estate scams, lottery scams, and [payroll fraud](#).

Why make the distinction between BEC and BAC? It's important when we talk about initial access, ways to bypass MFA, and remediation. Did a threat actor have access to email? Or are all applications enrolled in Okta? These details matter.

Over time, we expect the lines to continue to blur between BEC and BAC.

In fact, we're seeing this already in the case of M365, where access to an email account is no longer the destination, but merely a jumping-off point or potential pathway to Workday.

So how bad is the current situation? In short: it's bad.

According to the [FBI's most recent Internet Crime Complaint Center \(IC3\) report](#), BEC attacks threatened to cause an additional half a billion dollars in losses in 2021, bringing the total to \$2.4 billion. When the next report comes out this year, we expect that figure to grow.

In the year ahead, we think every CISO and CFO should prioritize identity threats as these attacks grow in reach and financial impact.

Some quick background on our reporting and analysis process

We score any compromise of an employee's email account by an external threat actor as a BEC attempt. Our data includes incidents where a hacker was able to compromise the password of a user but MFA or conditional access policies (CAPs) prevented unauthorized access. While the threat actor didn't gain access to the account *at the time*, the password was compromised. In fact, our SOC saw threat actors bypass CAPs by simply turning on a VPN service or repeating the authentication process using a legacy protocol.

For brevity and clarity, we're also going to talk about compromise of cloud access identity provider accounts, like Okta or OneLogin, under the identity threat umbrella, but we'll call it BAC. (See what we did there?)



TL;DR

WHAT WE SAW LAST YEAR

BEC accounted for 50% of all incidents, with 99% occurring in M365. Eleven percent of BEC attempts in M365 bypassed MFA using legacy protocols. AiTM phishing attacks to gain initial access and bypass MFA are on the rise.

- Representing 50% of all incidents, BEC remained the top threat to our customers, consistent with 2021.
- More than 99% of BEC attempts were in M365; fewer than 1% occurred in Google Workspace.
- Fifty-three percent of organizations experienced at least one BEC attempt; one was targeted 104 times.
- Six percent of BAC attempts used push notification fatigue to satisfy MFA.
 - This happens when an attacker continuously sends push notifications until the employee "authorizes" or "accepts" the request, allowing the attacker to satisfy MFA.
- Our SOC identified many BEC attempts targeting human capital management systems—specifically, Workday.
 - Once threat actors access Workday, they modify a compromised user's payroll settings to add the attacker's direct deposit information, thereby depositing the victim's paycheck into the attacker's account.

While some threat actors might want access to your email for fraud purposes, others have found that the data behind your applications also provides a pathway to perpetuate fraud.

More and more organizations use cloud access identity providers like Okta or OneLogin to provide a single sign-on (SSO) experience for their employees. This means an attacker can use a stolen credential to access more than just email.

What we saw last year: BEC and BAC identity-related threats top the charts

Representing 50% of all incidents, BEC attempts remained the top threat to our customers, which is the same when compared to 2021.


Of the BEC attempts we identified: more than 99% occurred in M365 and less than 1% were in Google Workspace.

Sixty percent 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 Identity and OneLogin.

Chart 1 shows the percentage of identity attack attempts where a threat actor targeted an email account compared to an account associated with a cloud identity provider.

Many of our customers use SSO, specifically Security Assertion Markup Language (SAML) authentication. This means that attackers can hit more than just email providers as easy targets to harvest credentials.

Chart 1: BEC (email compromise) vs. BAC (application compromise) in 2022




BEC targeting by industry

Threat actors targeted non-profit organizations most, followed by retail.

Chart 2 shows the percentage of BEC attempts identified by our SOC and demonstrates that BEC fraud isn't an industry-specific problem. A company's yearly revenue is by no means a predictable measure of potential BEC targeting either—BEC attempts to perpetuate fraud can happen anywhere, to any organization.

Representing 50% of all incidents, **BEC attempts remained the top threat to our customers**, which is the same when compared to 2021.

Chart 2: Top 10 targeted Industries




BEC targeting by business unit

Threat actors targeted employees working in sales departments the most, followed by operations and administration.

Chart 3 shows the percentage of BEC attempts by business unit identified by our SOC. Why the apparent targeting of sales roles? Because they focus on building relationships *outside* the organization. That means employees in these roles likely open a lot of emails from external addresses.

The takeaway here? If you invest in training so employees learn to recognize potential red flags associated with phishing emails (spoiler alert: you should), consider spending extra time with your sales organizations.

Chart 3: BEC targeting by business unit



BEC targeting by role

BEC threat actors don't just target CEOs or CFOs. Everyone in an organization is a potential target.


Hackers targeted employees working in individual contributor roles in 68% of BEC attempts. Manager, director, and executive titles accounted for the remaining 32%.

Chart 4 shows targeting by employee role. Our data shows that BEC targeting is likely a bit of a numbers game, given there are far more employees working in individual contributor roles than in management.

If you **invest in training so employees learn to recognize potential red flags** associated with phishing emails (spoiler alert: you should), consider spending extra time with your sales organizations.

Even though director-level and executive-level targeting accounted for only 19% of BEC attempts, it's worth noting that employees at these levels *are* targeted. Security controls must span the entire organization.

Chart 4: BEC targeting by role



BEC targeting frequency

BEC incident volume represented as a total percentage is one important part of the story, but BEC targeting frequency helps create a clearer picture.

Of the organizations targeted by BEC threat actors:

- Fifty-three percent experienced at least one BEC attempt.
- Twenty-four percent experienced at least three BEC attempts.
- One organization was targeted 104 times in 2022.

The takeaway? BEC happens frequently and remains a constant threat to organizations.

Significant BEC findings

Attackers abandoned legacy protocols for AiTM phishing

Threat actors have started moving away from authenticating using legacy protocols to bypass MFA in M365. Instead, they've adopted the use of frameworks such as [Evilginx2](#) to steal login credentials and session cookies for initial access and subsequent bypassing of MFA.

The net impact?


BEC threat actors successfully gaining access to M365 accounts is trending up. A main culprit? AiTM phishing to steal session cookies. Time-based one-time passwords (TOTP) and push notifications for MFA help protect against credential stuffing and password spraying, but they don't stop AiTM phishing and session cookie theft.

Yes, 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 TOTP and push notifications for MFA. While these factors certainly raise the bar considerably, gaps remain and AiTM attacks will continue to exploit those gaps.

Chart 5 shows the trend of successful account compromise over the course of 2022. This graph shows whether a preventive control stopped an attacker from accessing an account using a compromised credential.

Our data showed that as AiTM attacks increased, the success rate to access compromised accounts followed.

Chart 5: BEC success rate to accessing compromised accounts



Microsoft 365 session cookie theft is on the rise

One of the most notable findings in our data: session cookie theft for initial access and to bypass MFA is on the rise.

AiTM phishing sites harvest passwords and hijack a user’s sign-in session, bypassing the authentication sequence even if MFA is enabled.

Threat actors can use frameworks such as Evilginx2 to steal login credentials and session cookies.

Before we get into the details, it’s worth noting that FIDO2 keys and certificate-based authentication stop AiTM attacks.

Here’s how a typical AiTM attack plays out:

- An employee receives a phishing email with a malicious link to an attacker-controlled website, which acts as a proxy for the genuine website.
- The employee submits their credentials and, after prompting, provides an MFA challenge answer.
- The attacker proxy relays the requests and responses back and forth between the employee and the genuine site.
- The MFA challenge is authenticated by the employee and the website confirms the session is valid; a session token is then created and returned in the form of a cookie.
- The attacker-controlled proxy intercepts the session cookie and redirects the employee to the application URL the cookie was set up with.
- The attacker imports the session token cookies into their web browser and authenticates and bypasses MFA.

Once hackers access the email account, they typically query the email inbox for the phishing email that contains a link to their proxy site. Then they move the email to the deleted items folder to hide evidence of the attack. Finally, they register a new MFA device to establish persistence before the session cookie expires (as seen in Image 1).

Image 1: Alert for a new MFA device registered from an atypical location

(7) Azure MFA Device Added

MFA device added from IP with bad reputation

Kibana | Tune | Rule | Customer Context | GUIDs | RAW | Prosecutor | Report Issue | Pivot to...

Detection

Name

Azure MFA Device Added

Description

Identifies when a new MFA device is added

Supported Tech

Office 365

Mitre Tactics

Persistence

Details

Expel alert time

Vendor alert time

Vendor

Expel alert name

Description

Message

Office 365 • Office 365 Direct

Expel - Authentication Method Added

The user account

added an authentication method to the user account

United States

IP Flow Evidence

Source IP

Authentication using legacy protocols to bypass MFA is decreasing

Eleven percent of BEC attempts in M365 bypassed MFA using legacy protocols.

How does this work?

Original deployments of M365 tenants enabled IMAP and POP3 in M365 Exchange (as well as BasicAuthentication) by default.

IMAP and POP3 don't support MFA, so even if you have MFA enabled, attackers can still access these mailboxes.

BasicAuthentication allows attackers to authenticate with clients past any pre-authentication checks to the identity provider, which can lead to unwanted account compromises or account lockouts from password spray or brute force attacks.

What can you do? *Disable legacy protocols like IMAP and POP3 immediately.* This step is critical, especially if you've gone through the process of enabling MFA. Once you turn them off, strongly consider disabling BasicAuthentication to prevent any pre-auth headaches on your M365 tenants.

There is some good news.


Starting on October 1, 2022, Microsoft began disabling BasicAuthentication for Exchange Online. Chart 6 shows how the success rate of bypassing MFA using a legacy protocol has plummeted. Between Q1 2022 and Q3 2022, the success rate increased. Starting in Q4 2022, when Microsoft began disabling BasicAuthentication for Exchange online, we saw a significant decrease.

By December 2022, attempts to bypass MFA using legacy protocols dropped to zero—that’s right, zero.

Why didn’t the number immediately drop to zero once Microsoft began rolling out the change to disable BasicAuthentication?

Rolling out the changes *started* in October, but effectively implementing a change on such a widely used critical service can take years. In early January 2023, Microsoft announced that [time is up](#) and the exceptions they had granted were over—they turned off BasicAuthentication for Exchange Online tenants.

Chart 6: Trend of success rate to bypass MFA in M365 using legacy protocols



Threat actors targeted Workday to perpetuate payroll fraud

Starting in July 2022, our SOC identified identity attack attempts across multiple customer environments, targeting access to human capital management systems—specifically, Workday.

The goal of these attacks? Payroll and direct deposit fraud.

How hackers get into Workday

An attacker begins by compromising a user’s M365 or Okta account, often using BasicAuthentication (BAV2ROPC, IMAP, POP3) or stolen session cookies to bypass MFA—usually originating with VPN and hosting IPs.

From there, attackers access the victim’s Workday account directly through Okta, the compromised password, or a password reset email.

In scenarios where attackers compromise an M365 account and don’t have direct access to Workday via SSO, they read available documentation on payroll systems and new employee payroll enrollment. In most cases, they strive to gain access to human capital management systems using new employee setup procedures or password reset requests.

(Side note: we’ve also seen cases where attackers don’t use BasicAuthentication and the compromised user authorizes an MFA notification for the attacker using brute force Duo push requests. This involves an attacker continuously sending Duo push notifications to victims until they accept.)

Attackers can then create inbox rules within the compromised user’s email account to delete or move emails related to workday.com, myworkday.com, and/or emails with keywords like “payroll” or “assistance needed.”

To prolong this access, attackers can enroll trusted devices through an organization's mobile or endpoint device management platform (for example, Microsoft Intune).

Now, the attacker can modify the compromised user's settings to add the hacker's direct deposit information, depositing the victim's paycheck into the attacker's account.


IP address geolocation: half are geolocated in the United States

A notable finding in our data showed that nearly half of all authentication requests into a compromised email account originated from an IP address geolocated in the United States. In fact, the percentage increased throughout the year. While the majority of those IP addresses are associated with popular VPN services, we occasionally observed the use of IP addresses associated with residential internet service providers (ISPs).

Chart 7 shows the percentage of logins into a compromised email account originating from an IP address geolocated in the US or elsewhere each month. The percentage of logins originating from IP addresses in the US, here plotted in light blue, increased over the course of the year.

While the technique of acquiring and using residential proxies to perform BEC attacks isn't new, our data shows that attackers are using the technique more frequently, potentially allowing a BEC threat actor to bypass Okta network zones and conditional access policies.

Chart 7: Percentage of authentications into compromised accounts originating from US-based IP addresses



OAuth apps used to bypass MFA

Another notable finding in our data showed that less than 1% of the BEC attempts in M365 bypassed MFA by abusing OAuth applications. In this scenario, a BEC threat actor used phishing to trick the victim into granting permissions to an OAuth application. When the victim granted permission to the application, a security token associated with the victim was sent to the BEC threat actor (the application developer), which allowed access to the victim's data.

Typical remediation steps (clearing sessions, resetting the victim's password and MFA token) won't work in this scenario. Security or IT teams must remove the OAuth application and revoke its permissions. Restricting users from registering

new applications to Azure Active Directory (AD) can reduce the risks associated with malicious OAuth applications.

Significant BAC findings

Anomalous MFA push notifications

Our SOC found that 23% of the BAC attempts originated from suspicious MFA push notification activity. We've built a lot of detection muscle around spotting atypical Duo push notifications.

Duo authentication logs will show two devices and their IP addresses when a user (or threat actor) initiates a Duo push authentication:

- Access device—the device from which the user is requesting the Duo push
- Authentication device—the device approving or rejecting the Duo push

For a legitimate authentication event, we assume the access device and the authentication device are in close proximity, and since we have the IP addresses of both devices, we alert when there's an anomaly.

Employee-reported suspicious activity

We detected 40% of BAC attempts in Okta through employee-reported suspicious activity. In this context, an employee reported a notification of potentially suspicious activity from Okta or Duo. The lesson here? Have a defined, repeatable process to investigate and encourage your employees to report questionable events.

Push notification fatigue attacks to satisfy MFA

Six percent of BAC attempts used push notification fatigue to satisfy MFA. Push notification fatigue occurs when an attacker continuously sends push notifications until the employee “authorizes” or “accepts” the request, allowing the attacker to satisfy the MFA challenge. Again, encourage your employees to report activity like this.

How to spot BAC attacks

What can you do to detect—and hopefully prevent—these costly attacks?

Here's what we recommend for security teams:

- Alert for new Outlook inbox rules created with suspicious names (two to three characters in length, or repeating characters could be a clue). Also watch out for certain keywords, like “payroll” and “Workday.”
- Alert for multiple Okta sessions from the same user with multiple, non-mobile operating systems.
- Alert for potential brute force Duo push requests.
- Review any authentication using legacy protocol (UA = Bav2ropc) into M365, as it may represent MFA bypass. (Have you disabled legacy protocols yet?)
- Advise employees to maintain awareness of their information in the payroll system and report any abnormal or suspicious activity for further review. Investigate any unexplained variance in paychecks.

Have a defined, repeatable process to investigate and encourage your employees to report questionable events.



How to protect your organization from BAC

- Deploy phish-resistant MFA. (FIDO security keys for the win!)
- If FIDO-only factors for MFA are unrealistic, disable email, SMS, voice, and TOTPs. Instead, opt for push notifications.
 - Then configure MFA or identity provider policies to restrict access to managed devices as an added layer of security.
- Implement a pre-auth policy for network zones in Okta.
- Consider blocking access to Okta from suspicious network zones based on IP address(es), autonomous system numbers (ASN), IP types, or geolocation.
 - Why? Clients from blocked zones can't access Okta URLs and requests are automatically blocked before authentication.
- Deploy Okta's adaptive multi-factor authentication (AMFA).
 - It reduces risk by blocking authentication attempts with previously unseen authentication characteristics, such as impossible travel, unusual locations for the environment, or a new device for the account. Admins can define the actions Okta takes and the variables it considers through policies in the Okta console.

Want to learn more about how
Expel can help with **identity threats**?

[Let's Chat](#)



Cloud platform security trends

Attackers go after all major platforms using a mix of tactics...

We continue to see attackers gaining unauthorized access or compromising cloud resources in Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure. We see both opportunistic and targeted attacks, with exposed keys, stolen user credentials, or public-facing resources leading to initial access.

For context, the end goals for attackers range from deploying commercially available malware and cryptocurrency miners to cloud-hosted servers in the data plane—where cloud resources reside—to attempting to gain administrator access to the control plane to establish persistence and deploy new resources through the cloud provider's web console or API calls.

What we saw: cloud attacks in 2022

Incidents identified by our SOC in the three major public clouds (AWS, GCP, and Azure) increased by 70% over 2021—underscoring the importance of an effective threat detection and response strategy in the cloud.

Of the incidents spotted by our SOC in one of the clouds:

- Ninety-two percent occurred in AWS
- Four percent occurred in Azure
- Four percent occurred in GCP

For context, we monitor roughly four times as many AWS environments as Azure and ten times as many as GCP.

In terms of patterns and trends, two major themes 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 or domain name system (DNS) rebind attacks in AWS, or to deploy tooling and malware.

Red team exercises or penetration tests accounted for 50% of cloud infrastructure incidents we spotted in 2022. We removed the red team data for the sections that follow, but we'll come back to these later.

The data in this section tilts heavily toward AWS, given the low volume of incidents we spotted in Azure and GCP.



TL;DR

WHAT WE SAW LAST YEAR

Cloud incidents increased 70% compared to 2021. Of the cloud security incidents we responded to, 20% resulted in deployment of cryptocurrency miners and 13% resulted in deployment of publicly available commodity malware.


- Threat actors started moving away from authenticating via legacy protocols to bypass MFA in M365. Instead, they adopted frameworks such as Evilginx2, facilitating AiTM phishing attacks to steal login credentials and session cookies for initial access and MFA bypass.
- FIDO2 (Fast ID Online 2) and certificate-based authentication stop AiTM attacks. However, many organizations don't use FIDO factors for MFA.
- Attackers targeted long-term access keys and service account credentials to gain initial access to public cloud environments.

Cloud attack vectors

Attackers accessed cloud infrastructure in one of two ways: they either compromised cloud credentials (including exposed long-lived credentials) or exploited a service or application running on a cloud compute instance (we call this configuration abuse).

Of the attackers who gained access to cloud infrastructure, 56% did so through cloud configuration or application abuse and 44% used compromised cloud credentials.

Chart 8: How attackers gained access to cloud infrastructure



From a defender's perspective, it's useful to think about the attack surface in the cloud on two planes: the "cloud control plane" and the "data plane." The cloud control plane owns the provisioning and configuration of cloud resources, while the data plane is where resources functionally exist and operate together. Why is this important? Each plane comes with its own risks, sources of evidence, and implications for incident response.

For example, imagine an attacker compromises a publicly facing Remote Desktop Protocol (RDP) server running on a cloud compute instance. This incident has occurred on the **data plane**. As a defender, you'll want to examine activity on that RDP server as well as network activity on its connected networks to look for lateral movement. We look at event logs, endpoint detection and response (EDR) activity, and virtual private cloud (VPC) flow logs to investigate and respond.

Now let's say we're dealing with compromised identity and access management (IAM) keys. This is an example of an incident occurring on the **cloud control plane**. As a defender, we'll want to look at the cloud activity logs (whether that's CloudTrail or an equivalent) to understand how those credentials were used by the attacker.

Finally, it's important to understand that attackers can (and will) traverse these planes during the course of an attack. For example, an attacker on a cloud compute instance could pivot to the cloud control plane by stealing instance-assigned IAM credentials. Alternatively, an attacker with privileged IAM credentials could create their own instance to pivot to the data plane for further compromise.

Let's dive deeper into compromised cloud credentials and what we saw on the cloud control plane.

Incidents identified by our SOC in the three major public clouds (AWS, GCP, and Azure) **increased by 70% over 2021**

Compromised cloud credentials

Attackers targeted long-term access keys and service account credentials as a means for initial access.


Of the cloud credentials compromised by attackers:

- Seventy-five percent were through long-term access keys and service accounts.
- Twenty-five percent were through cloud console credentials.

Why do long-term access keys and service accounts make up 75% of all compromised cloud credentials?

Long-term credentials are used by applications and resources to programmatically interact with cloud resources. As a result, they're more likely to be accidentally exposed in code repositories and version control software. Successful server-side request forgery and DNS rebind attacks may also expose AWS long-term access keys from the metadata service. User credentials will likely require social engineering to acquire them.

Chart 9: Cloud credential compromise by category



So what happens once an attacker gains entry to the cloud control plane?

For starters, they typically interact with the control plane using software developer kits (SDKs) like Boto3 and Google Cloud SDK as well as command line utilities like AWS CLI and PowerShell.

In 83% of these incidents, the attackers used an SDK or CLI, and authenticated from proxy or hosting infrastructure. Seventy-five percent of these were from outside of the US.

Image 2: Alert view in Workbench™ for an SSH key pair that was created using the AWS CLI by an attacker

Detection

Name

Description

Supported Tech

Mitre Tactics

Suspicious SSH key pair generated

Identifies suspicious generation of AWS SSH key pairs which may represent account backdooring.

AWS CloudTrail

Credential Access

Details

Expel alert time

Vendor alert time

Vendor

Source ID

Vendor alert name

Description

2022-11-07T13:53:09Z (2 months ago)

2022-11-07T13:45:53Z (2 months ago)

 .

Expel - AWS SSH key pair generated

User type - IAMUser MFA Used -

IP Flow Evidence

Source IP



User Evidence

Username

SID

Email address

URL Evidence

User agent

aws-cli/1.24.10 Python/3.6.9 Linux/5.4.0-1085-aws botocore/1.16.12

We can use this data to see that attackers were more likely to leverage proxy infrastructure with long-term credentials based outside of the U.S. Although we can't predict if this trend will continue, it's likely a good starting point for detecting and monitoring long-term credentials used in your environment on some of the APIs we'll call out below.


Once attackers have gained access to a cloud environment, they'll likely follow a series of standard steps. We typically see a series of API calls to enumerate the AWS environment and establish some type of persistent access to continue post-exploitation.

Here's a breakdown of common attacker tactics we see in AWS incidents:

- Sixty-seven percent issued reconnaissance and discovery API calls.
- Forty-two percent established persistence by creating new IAM users, access keys, and key pairs.
- Twenty-five percent performed defense evasion by deleting credentials and IAM roles.
- Twenty-five percent impaired defenses by modifying and creating firewall and routing rules.
- Seventeen percent attempted to execute commands using APIs such as ExecuteCommand.

In our [AWS](#) and [GCP](#) mind maps, we call out APIs that attackers use or could use with unauthorized access to cloud environments. These investigative resources are key for building our detection methodologies. Chart 10 shows the specific breakdown of API calls most commonly used by attackers in AWS this past year.

Chart 10: Visual representation of APIs most commonly used by attackers as seen in incidents




Cloud configuration abuse and the data plane

Attackers frequently exploited or abused misconfigurations or applications hosted on cloud resources as a means of gaining initial access on the data plane or interacting with the AWS metadata service. Fifty-six percent of our attacker-related incidents occurred as a result of configuration abuse.

Here’s a breakdown of the incidents spotted by our SOC on the cloud data plane:

- Twenty-seven percent were the result of the deployment and installation of cryptocurrency miners.
- Twenty percent were server-side request forgery attacks to acquire AWS metadata.
- Twenty percent were DNS metadata rebind attacks to query AWS metadata.
- Twenty percent were the result of deployment of attacker utilities and attacker activity.
- Thirteen percent were from the deployment and installation of publicly available commodity malware.

Chart 11: Breakdown of configuration abuse on the cloud data plane



Red team exercises and penetration tests in the cloud

Red team exercises and penetration tests give us a unique opportunity to detect and investigate post-exploitation activity that isn't remediated in a customer environment. With 50% of incidents in cloud infrastructure detected as authorized red teaming or penetration testing, we wanted to share what we saw.

Of note: we distinguish between red teams and penetration tests, as red teams orient more toward emulating adversarial behavior, with clear objectives to identify risk. Pen tests gear more toward auditing the environment to identify risk.

Here's what we found:

- Red team exercises and penetration tests emphasized enumerating AWS environments to review security controls, policies, and configurations within AWS services. Organizations provided red teams and penetration testers with an IAM user or limited permissions credentials prior to the exercise.
- Penetration tests placed less emphasis on obfuscating their source IP addresses and User-Agent strings. Fifty percent of pen tests used cloud auditing tools or Kali Linux to interact with AWS and 85% used legitimate ISPs to authenticate to the cloud.
- Red teams emphasized post-exploitation activity, such as attaching administrator roles, creating and starting resources, sharing snapshots, and authenticating with externally managed users.
- Red teams and penetration testers in the data plane focused on exploitation, asset discovery, and network service scanning across pods and EC2 instances.

Other risks in the cloud

Not all risk in the cloud relates to compromise or attacker activity. Business risk for cloud infrastructure includes resources made public, resources deleted or created, and overuse of root or service account credentials. To reduce risk we deliver business risk notifications to our customers via our orchestration bot, Ruxie™, once they determine it's something they need to hear about.

Based on their configurations, we delivered these notifications to our customers:

- AWS root authentications (40%)
- Deleted GCP KMS keys (40%)
- S3 bucket set to public (20%)

Cloud compromise detection opportunities

- For AWS, create a detection methodology around long-term access keys (these keys start with "AKIA") and short-term access keys (these keys start with "ASIA") to detect the early stages of unauthorized access to your environment.
- Alert when you see long-term access keys used from unauthorized regions or ASNs that are VPNs, hosting providers, or data centers.
- 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.
- Alert on new access keys and key pairs created from anomalous regions or abnormal IP addresses.

50% of pen tests used cloud auditing tools or Kali Linux to interact with AWS and 85% used legitimate ISPs to authenticate to the cloud.

- Trigger an alert for a high number of failed 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 or infrastructure types, and device types.



How to protect your organization

- Separate out high-privilege permissions required to perform work into different groups and roles, and only provide access to those permissions as needed to different users.
- Avoid use of static IAM credentials where possible in favor of short-lived, just-in-time access credentials (like IAM roles in AWS or service account impersonation in GCP).
- Periodically rotate permanent access keys. Pro tip: if you're accidentally exposed, scheduled rotations limit the timeframe an attacker can use access keys to perform unauthorized actions in your environment.
- Limit root access and usage.
- Enable IMDSv2, the second version of the instance metadata service, and prevent unauthorized requests to it.
- Enforce least privilege controls for IAM resources and users.
- Enforce MFA for access to cloud consoles. Pro tip: these days, most cloud infrastructure providers have robust SSO integrations that seamlessly provide short-lived credentials for console or API access.

Want to learn more about how
Expel can help with **cloud security**?

[Let's Chat](#)



Ransomware

Ransomware is a top priority for most organizations we protect. The proportion of organizations [attacked by targeted ransomware threat groups](#) in 2022 doubled from 2021.

But before we get into the data, it's important to provide some background. As a security operations provider offering managed detection and response (MDR), our goal is to stop attacks before they result in a ransomware incident. You don't need detections to tell that your files have been encrypted—the ransomware does that for you. We work to make sure it doesn't get to that point. The data we share focuses on how and when in the attack lifecycle we identified and stopped ransomware attacks *before* an encryption event. We focus on detecting the initial compromise or spotting when a ransomware threat group establishes a foothold within a target organization.

What we saw last year: more pre-ransomware attacks

Eleven percent of incidents could have resulted in deployment of ransomware had we not been there to intervene—a seven percentage point increase over 2021.


Our data shows that when our SOC responded to an incident deployment or installation of malware, there was a 32% probability the activity was related to ransomware operations.

Targeting by industry

Among the industries we serve, the manufacturing industry experienced the greatest proportion of ransomware incidents. The entertainment and legal services industries followed, each accounting for 12%.

But, like identity attacks, the data also tells us that ransomware attacks aren't an industry-specific problem. Neither industry nor annual revenue provide predictable measures of potential ransomware targeting. This is an everywhere-and-everyone problem.

Chart 12: Ransomware targeting by industry



TL;DR

WHAT WE SAW LAST YEAR

Pre-ransomware activity accounted for 11% of all incidents—an increase of seven percentage points vs. 2021.

Attackers used zipped JavaScript files to gain initial access in 44% of all ransomware incidents, while they used ISO files to gain initial access in another 12%. Something interesting: we didn't see attackers using ISO files at all in 2021.


- Microsoft is making it easier to block macros in files downloaded from the internet, so ransomware threat groups and their affiliates are abandoning their use of VBA and Excel 4.0 macros to gain initial entry to Windows-based environments. Instead, they're opting for ISO, LNK, and HTA files.
- Nine percent of all ransomware incidents started from an infected USB drive.

Ransomware attack vectors

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 fourteen percentage points compared to 2021).
 - For reference, “self-installation” occurs when employees unintentionally install malware by double-clicking a zipped executable or enabling a malicious macro in an Office doc. These are not exploits.
- 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 13: Top ransomware attack vectors by category



Ninety-seven-and-a-half percent of all ransomware incidents we identified used a self-installation attack technique to gain initial entry to a computer running Windows. Here’s a breakdown of those techniques:

1. Zipped JavaScript file (44%)
2. Excel macros (16%)
3. ISO file (12%)
4. Removable media (9%)

Chart 14: Top ransomware attack vectors



Eleven percent of incidents could have resulted in deployment of ransomware had we not been there—to intervene—a seven percentage point increase over 2021.

Something to note: Our data showed that attackers used ISO files to gain initial access in 12% of all ransomware incidents. This attack vector didn't make our list in 2021.

So why such a significant increase in the use of ISO files for initial access?

As Microsoft continues to make it easier [to block macros in files downloaded from the internet](#), ransomware threat groups and their affiliates have started to abandon their use of VBA macros and Excel 4.0 macros to gain initial entry to Windows-based environments. Instead, ransomware operators opted to use ISO, LNK, and HTA files.

As we evolve as defenders, so too do the attackers—unfortunately.

And don't forget about removable media! Infected removable media, especially when inserted and connected to an operating system with "autorun" capabilities, can cause a real headache. Nine percent of all ransomware incidents started from an infected USB drive.

What's more, when responding to pre-ransomware activity that started from an infected USB drive, we'd likely see a variant of the Raspberry Robin family of malware deployed. This worm will sell access to a compromised network to [ransomware threat groups and their affiliates](#), and the family of malware is also known to deploy Cobalt Strike's Beacon agent to infected computers.

Speaking of Cobalt Strike's Beacon agent, we observed it in 3% of all incidents in 2021. In 2022, that number decreased to fewer than 1% of all incidents.

Why the drop? Our data shows that we stopped pre-ransomware activity earlier in the attack lifecycle. Of the ransomware incidents identified by our SOC, only 3% of incidents were spotted and remediated after the initial access phase. Attackers typically deploy Cobalt Strike's Beacon agent later in the attack lifecycle, once they've established initial access and elevated privileges.

In summary, monitoring and reducing the self-installation attack surface on the Windows OS by associating Windows Script Host (WSH) files with Notepad and disabling XLM and VBA macros in files downloaded from the internet are key to defending against ransomware attacks.

But what about those pesky ISO files?

Block ISO files at email and web gateways. But proceed with caution, as many businesses use these files in the regular course of business. Our recommendation? Consider unregistering ISO file extensions in Microsoft Windows Explorer. By doing so, Windows will no longer recognize ISO files and double-clicking won't result in program execution.

Top malware families

As stated above, only 3% of the ransomware incidents we identified made it past the initial infection phase before we stopped them. In other words, we stopped a potential ransomware incident before ransomware could be delivered 97% of the time. Our SOC typically stopped the incident either immediately after an attacker exploited a web server or a malicious downloader, or after a stager executed on a workstation. When analyzing the malware families that attackers deploy for ransomware events, we include the incidents that we stopped before ransomware delivery.

Nine percent of all ransomware incidents **started from an infected USB drive.**


The top malware families we attributed to identified ransomware operations were:

1. Gootkit loader (12% of all incidents)
2. Qakbot (11%)
3. SocGholish (5%)

Initial lead by class of technology

Chart 15 shows the initial leads into ransomware activity from our [various technology integrations](#). While a significant percentage of initial leads into pre-ransomware activity originated from EDR devices, it's worth noting that integrations with various network detection and response (NDR) and security information and event management (SIEM) technologies also enabled our SOC.

Chart 15: Initial lead source into ransomware activity by technology class



Top Windows processes to monitor

Attackers favored the genuine Windows-based script host process, wscript.exe, as their top process executed for initial access. What does this mean? Organizations should monitor native Windows OS binaries and popular M365 productivity applications for signs of malicious activity related to threat actors "living off the land."


If you're looking for more general guidance on which Windows processes you should monitor, our ransomware incident data indicates that the following were the top processes attackers used to gain entry in 2022:

1. WScript.exe (29% of all incidents)
2. CMD shell (16%)
3. PowerShell (14%)
4. RegSvr32.exe (5%)
5. Rundll32.exe (5%)

We stopped a potential ransomware incident before ransomware could be delivered 97% of the time.

Chart 16 shows the percentage of Windows processes executed to gain initial access to a network attributed to ransomware operations.

Chart 16: The top 10 Windows processes executed to gain initial access



Top hacker behavioral patterns

Our SOC recorded some of the top behavioral patterns that indicate pre-ransomware activity, and it did so before any damage could be done. Alerting on behavioral patterns like these helps us spot all kinds of malicious activity, not just ransomware threat groups trying to infiltrate the network. However, since every environment differs and typical user behavior can differ by organization, organizations may need to conduct some tuning and filtering to keep these patterns from alerting on legitimate behavior in your environment.

Endpoint behavioral patterns alert when...

- Installation:
 - ▢ A scripting process other than PowerShell (like wscript.exe) launches a PowerShell process with encoded commands.
 - ▢ The Microsoft Mshta utility, mshta.exe, loads a command line from an entry stored within the Windows registry.
 - ▢ The Microsoft Excel process, Excel.exe, launches the Register Server (regsvr32) to execute a file from the active Windows User profile.
 - ▢ Suspicious behaviors related to scripting processes like wscript.exe or cscript.exe:
 - Execute a .vbs, .vbscript, or .js file from a Windows user profile.
 - Initiate an external network connection.
 - Spawn a cmd.exe process.
 - ▢ Command and control (C2): PowerShell executes a base64 encoded command and the process initiates an external network connection.

Network behavioral patterns alert when...

- Installation: A PowerShell network connection to a rare external destination appears.
- C2: An HTTP connection with a PowerShell user-agent appears.
- Elevate: 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 HTTP-based outbound connections with a PowerShell user-agent?” 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.



How to protect your organization

Our top five recs to protect your organization from ransomware attacks:

1. 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.
2. 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.
3. Don’t expose remote desktop protocol (or any other service you don’t need to) directly to the internet.
4. Disable macros in files downloaded from the internet. Microsoft makes this setting a default, but it’s important to know if they’re still enabled for your organization.
5. Patch and update regularly. (Just do it!)

One last piece of advice: have a ransomware incident response (IR) plan and test it. A real-life security incident isn’t the best time to test your IR plan. A great IR plan 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.

Want to learn more about how
Expel can help with **ransomware**?

[Let’s Chat](#)



Phishing

Phishing attacks serve as an entry point for many attack types (including the deployment of commodity malware and ransomware—not just BEC), and they continue to significantly challenge organizations of all sizes.

Our phishing team collects suspicious emails from our customers to determine if they're malicious, unwanted spam, or legitimate communications. This gives us unique visibility into the various phishing attacks launched to gain initial access to a target organization, or even phish for [AWS account root user credentials](#).

What we saw from phishers last year: lots of credential harvesting

Sixteen percent of phishing email submissions 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 was composed of 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."

The top subject lines in malicious emails that resulted in an employee click or compromise were, "Incoming Voice Message," "Checking in," and "Voice Mail Call received for <user's email>." Our data shows that social engineering themes that are actionable, time-sensitive, and financially driven are most successful. (FYI, the data in Chart 17 is focused on all malicious attempts, the specific subject lines mentioned previously are those that successfully led a user to click a link and/or resulted in compromise.)

Targeting by industry

Among the industries we serve, the hospitality industry experienced the greatest proportion of phishing attacks. Finance and legal services followed, each accounting for 13% of attempts spotted by our phishing team.

Why? Our data showed organizations that frequently correspond with a high number of third-party vendors were most likely to experience a successful phishing attack. Also worth noting is that **these organizations were the most likely to perform a phishing simulation to help educate employees about these attacks and their risks.**



TL;DR


WHAT WE SAW LAST YEAR

Credential harvesters accounted for 88% of malicious emails.

The top subject lines in malicious emails that resulted in an employee click or compromise were, "Incoming Voice Message," "Checking in," and "Voice Mail Call received for <user's email>."

- Credential theft via phishing continues to grow, with identity the main focus of today's attacks.
- Our data shows that social engineering themes that are actionable, time-sensitive, and financially driven are most successful.

Chart 17: Successful phishing attempts by industry



Most used subject lines in phishing emails

The table below shows a sampling of subject line keywords we see regularly across our customer base. Our data shows that social engineering themes that create urgency, FOMO (fear of missing out), intrigue, or potential financial loss spur the most action (open, click, interact).

Keyword	% of total submissions	Example subject lines
Delivery	13.3%	- Delivery Confirmation - Fax Delivery Report - FedEx Delivery Notification
Parcel	6.8%	- Parcel Arrival Notice - Parcel Delivery - Parcel Redelivery Service Notice
Request	5.2%	- You've got a money request - Business Proposal Request - Request
Order	4.7%	- Order to appear and present on behalf of your company - Order Confirmation - Your Order Has Been Placed!
Payment	4.7%	[Health Payment] Encrypted Message - Aetna EFT Payment Notification - EFT Settlement Payment Approved
Invoice	4.3%	- Invoice No: <number> - Monthly Statement//Past Due Invoices - Payment Due Immediately - URGENT : INVOICES
Password	1.2%	- Password expiry - Password Expires Today - Password Warning
Urgent	0.9%	- Urgent request - URGENT : INVOICES - URGENT: Payment declined
Available	0.8%	- Available? - Are you available now? - Are you available at the moment?
Response	0.8%	- Response - Quick Response - Response Required

Table 1: Top ten subject line themes in malicious emails



How to protect your organization from phishing

1. Make sure you're running MFA wherever possible and using phish-resistant FIDO security keys to significantly reduce the risks associated with credential theft.
2. Consider deploying a secure email gateway (SEG) to monitor incoming and outgoing emails for signs of an attack.
3. Invest in training so employees learn to recognize potential red flags associated with phishing emails.
4. 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.
5. Use email anti-spoofing controls such as DMARC, SPF, and DKIM.

Want to learn more about how
Expel can help with **phishing**?

[Let's Chat](#)



Looking ahead to 2023...

Thoughts from our SOC leadership team

Dan Whalen

Manager, Research & Development:

As I think about the year ahead, I'm reminded that security doesn't exist in a vacuum. Looking around at the macro-economic situation, I believe businesses in general are looking to do more with less. Often, that means leaning on technology like cloud infrastructure, software-as-a-service, or even adopting tools that only a few years ago were viewed as radical and cutting-edge (looking at you, Kubernetes). Fast adoption of technology always comes with some risks. Unfortunately, organizations often learn this lesson the hard way (when they experience a security incident). As organizations adopt new technology to react to changing markets, I anticipate we'll see an uptick in security incidents resulting from misconfigurations (easy to miss, and easy for attackers to capitalize on).

Ray Pugh

Senior Manager, Security Operations:

Attackers love to prey on stressful situations. With economic uncertainty looming ahead in 2023, opportunistic threat actors will craft tailored campaigns to capitalize on the environment. Whether it's providing credentials or interacting with self-installed malware, the trends of user interaction leading to compromise are likely to persist as they remain the most accessible, effective means for gaining access to a target environment. Emphasis on end user training, backed by strong capabilities to detect and respond to threats that slip through the cracks are the best bet for defenders to keep pace as these threats evolve.

Ben Brigida

Director, Security Operations:

I predict that we'll see an uptick in MFA push notification fatigue attacks. Simply put, they work, and the fact that an increasing number of organizations turn to SSO to provide access to their enterprise applications means that they're juicy targets for attackers. Accessing SSO platforms means they get the "keys to the kingdom," so they're not shy about blitzing their targets with MFA push notification requests in the hopes that they'll eventually just approve one.

Luckily, security teams have a few options at their disposal for preventing this. First, they can disable the push notification feature of the MFA solution and require users to enter a secure PIN. Another option is to use the number-matching setting common in popular platforms, which requires the user to enter numbers from the platform to approve the authentication request.

Unfortunately, attackers will then try to find another way around these protections.

Steve Edwards

Director, Detection and Response:

In 2023, I predict that attackers will increase their focus on the implicit trust organizations have for their supply chain. This trend—which included newsworthy breaches over the past few years such as CircleCI, npm, and SolarWinds—continues to illustrate that we're only as strong as our weakest link.

To keep up, the industry will need to double down on restricting access to "least privilege" and "need to know" for these types of interconnections. They will also need to provide better insight into the activities occurring between application-to-application connections to allow for a more complete detection and response strategy to be developed. The good news is these data flows are predictable, which makes anomaly detection more simple to apply than in human-to-application interactions.

Jon Hencinski

Vice President, Security Operations:

Microsoft's disabling legacy protocols in M365 and macros in files downloaded from the internet has changed attacker behavior. But, with the advancement of good comes the advancement of evil.

My prediction? We'll see BEC threat actors continue to adopt the use of AiTM phishing frameworks to steal session cookies to gain initial access and bypass MFA. TOTP and push notification MFA do not stop this class of attack. As a result, BEC will become an even bigger problem in 2023. But perhaps this will be the compelling event that organizations need to shift to FIDO2.

About Expel

Expel is the leading managed detection and response (MDR) provider trusted by some of the world's most recognizable brands to expel their adversaries, minimize risk, and build security resilience. Expel's 24/7/365 coverage spans the widest breadth of attack surfaces, including cloud, with 100% transparency. We combine world-class security practitioners and our AI-driven platform, Expel Workbench™, to ingest billions of events monthly and still achieve a 23-minute critical alert MTTR. Expel augments existing programs to help customers maximize their security investments and focus on building trust—with their customers, partners, and employees. For more information, visit our [website](#), check out our [blog](#), or follow us on [LinkedIn](#).

WANT TO LEARN MORE?

- [Learn about the problems we solve](#)
- [Watch a video demo](#)
- [Take a tour of Expel MDR for Cloud Infrastructure](#)
- [Subscribe to our blog](#)
- [We're hiring! Find the right role for you](#)
- [See what Expelives say about working at Expel](#)

2022 at Expel



Awards

- [Cybersecurity Breakthrough Awards 2022: Threat Detection Solution Provider of the Year](#)
- [CISO Choice Awards: Security Operations Category](#)
- [2022 Deloitte Technology Fast 500™](#)
- [CRN® Emerging Security Vendors List](#)
- [CRN® Channel Chiefs List](#)



Industry Recognition

- [KuppingerCole's 2022 Market Compass for Security Operations Center-as-a-Service \(SOCaaS\)](#)
- [Frost & Sullivan: Global MDR Radar Report](#)



News

- [Expansion into EMEA](#)
- [Expel Investors Fuel Rapid Growth with Additional Series E Investment, Bringing Total Funding to \\$288.8 Million](#)
- [Total Economic Impact Study Reveals 610% Return on Investment for Expel Customers](#)



Product and Feature Launches

- [Expel Unveils Threat Research and Cloud Detection, Response and Remediation Capabilities and Resources](#)
- [Expel Showcases Impactful Product Advancements in Detection, Response and Measurement at RSA Conference 2022](#)



Great Place to Work

- [Great Place to Work® and Fortune Name Expel to Three 2022 Best Workplaces Lists](#)
- [America's Best Startup Employers](#)



Expel Quarterly Threat Reports

- [Q1 2022](#)
- [Q2 2022](#)
- [Q3 2022](#)

Appendix

Appendix A: alert-to-fix times

There’s nothing more strategic than defining where you want to get to and measuring your progress. Strategy informs what “good” means and measurements tell you if you’re there or not.

In building our SOC here at Expel, we set a target that for every incident we identify: we’ll respond faster than pizza delivery (30 minutes or less).

For 2022, our median alert-to-fix time across critical incidents was 28 minutes, and our fastest time was one minute and 35 seconds.

We define alert-to-fix as the elapsed time from when an Expel alert arrives in our queue to when the first remediation action is assigned to (or performed on behalf of) the customer.

Appendix B: SOC quality control


There’s a common assumption that you have to choose between scale and quality.

In building our SOC here at Expel, we didn’t want to trade quality for efficiency—it just didn’t feel right.

Everyday, we perform a quality control (QC) check based on a set of manufacturing standards to answer: “Is our work quality improving or are things getting worse?” You can find out more about our SOC QC program [here](#). The graph below shows trend data for errors across our categories of work performed in 2022.

The outcome? Our overall SOC quality improved even as we scaled—adding new team members, significantly expanding our customer base and managing more alerts and incidents.

SOC review 2022





Ready to talk to a human?

www.expel.com/contact-us

[844.397.3524](tel:844.397.3524)