expel

Quarterly Threat Report

Q3 2022

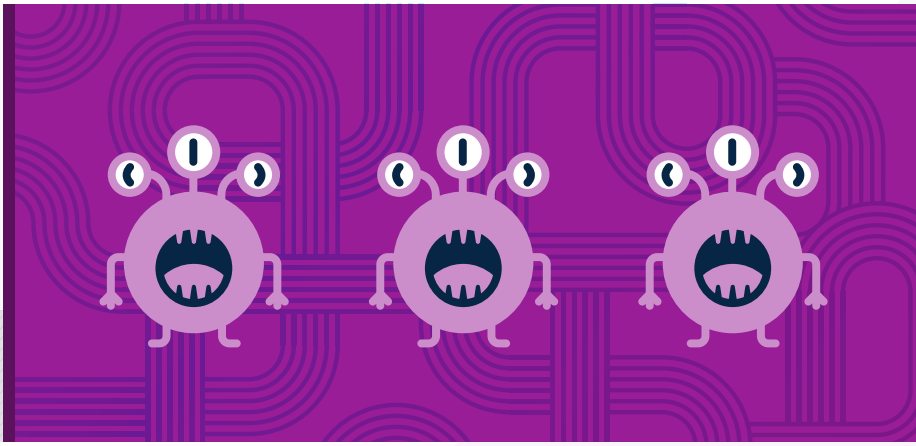

www.expel.com

Contents

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

Welcome to our third Expel Quarterly Threat Report. These reports 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.

These trends are based on incidents our security operations center (SOC) identified through investigations into alerts, email submissions, and threat hunting leads in the third quarter (Q3) of 2022. We analyzed incidents across our customer base, spanning organizations of various shapes, sizes, and industries, from July 1 to September 30, 2022. In the process, we distilled patterns and trends to help guide strategic decision-making and operational processes for your team. We used a combination of time series analysis, statistics, customer input, and analyst instinct to identify these key insights.

Our goal: by sharing how attackers got in and how we stopped them, we'll translate the events we detect into security strategy for your organization.

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

Identity is still the new endpoint, and it's showing no signs of slowing down.

- Identity-based attacks accounted for 59% of all incidents in Q3, an increase over Q2.
- Business email compromise (BEC) and business application compromise (BAC) accounted for 55% of all incidents identified.
- All BEC attacks targeted Microsoft 365. Note: that trend may change since Microsoft disabled Basic Authentication in Q4 by default.
- Effective detection and response strategies for more and more organizations will be identity-oriented.

Users increasingly let attackers in by approving MFA pushes for BAC.

- Only about half of the BAC incidents we spotted resulted in the attacker successfully gaining access to the account; the other half were stopped by multifactor authentication (MFA) or conditional access policies.
- MFA and conditional access were configured for over 80% of the successful compromises, but the attacker successfully tricked the legitimate user to satisfy the MFA request.
- To stop MFA push notification fatigue attacks, organizations can disable push notifications in favor of a PIN or a Fast Identity Online (FIDO) compliant solution.
- If that's unrealistic, control push notifications using number matching—a setting that requires the user to enter numbers from the identity platform into their MFA app to approve the authentication request.

Attackers use IPs geolocated in the U.S. when targeting U.S.-based organizations.

- Forty-eight percent (48%) of the BEC attempts and 46% of the incidents where an attacker successfully accessed an account originated from an IP address with a U.S. geolocation.

Attackers are more likely to **successfully bypass conditional access mitigation efforts** by using U.S. IP addresses.

- All of the authentication attempts originating from the U.S. came from an IP associated with a VPN or hosting provider.
- Attackers are more likely to successfully bypass conditional access mitigation efforts by using U.S. IP addresses.

Ransomware threat groups and their affiliates have turned to zipped Javascript or ISO files, continuing abandonment of the previously popular use of visual basic for application (VBA) macros and Excel 4.0 macros to gain initial entry to Windows-based environments.

- Zipped JavaScript files accounted for 46% of all pre-ransomware incidents and zipped ISO files another 26%.
- Remaining attacks employed removable media (10%) and a few Excel 4.0 macros (8%).

The top subject line themes for malicious emails were no subject line, followed by some iteration of, "Invoice," "Order confirmation," "Payment," and "Request."

- Our data shows social engineering themes that create urgency, a fear of missing out (FOMO), or potential financial loss are most likely to get a person's attention and result in action (open, click, interact).

MFA and conditional access were configured for **more than 80% of the successful compromises**, but the attacker successfully tricked the legitimate user to satisfy the MFA request.

Q3 by the numbers

Incident types

Identity-based attacks (credential theft, credential abuse, long-term access key theft) accounted for 59% of all incidents handled by our SOC in Q3—up three percentage points compared to Q2.

BEC (unauthorized access into email apps) and BAC (unauthorized access into application data) accounted for 55% of all incidents, an increase of four percentage points from Q2. Identity-based attacks in popular cloud environments like Amazon Web Services (AWS) decreased by two percentage points, accounting for 3%.

The deployment of commodity malware, cryptominers, and malware families linked to pre-ransomware operations accounted for 31% of incidents—down three percentage points from Q2.

As we get closer to the end of the year, we enter what Expel refers to as "red team season," the time of year when we see security teams working more with outside firms to conduct penetration testing. We began seeing the first instances of this: authorized penetration tests by red and purple teams accounted for 9% of the incidents our SOC detected.[1] This is roughly two percentage points more than we saw last quarter.

Activity associated with advanced persistent threats (APT) led to 2% of incidents, and they all came to us from incident response (IR) partnerships. APT groups are still active, but comprise a small percentage of total incident volume.

Looking at the bar graph below, we can't help but recognize how much of our jobs take place in a browser these days. On average, companies use 254 Software-as-a-Service (SaaS) applications. The number of SaaS apps organizations use continues to climb, which means more of the sensitive information we all need to do our jobs is protected by our credentials—not our (well-managed, secure, heavily monitored) workstations. Effective detection and response strategies for more organizations now are identity-oriented. Endpoint detection and response (EDR) tools don't provide broad enough coverage on their own.

Chart 1: Incidents detected by the Expel SOC in Q3

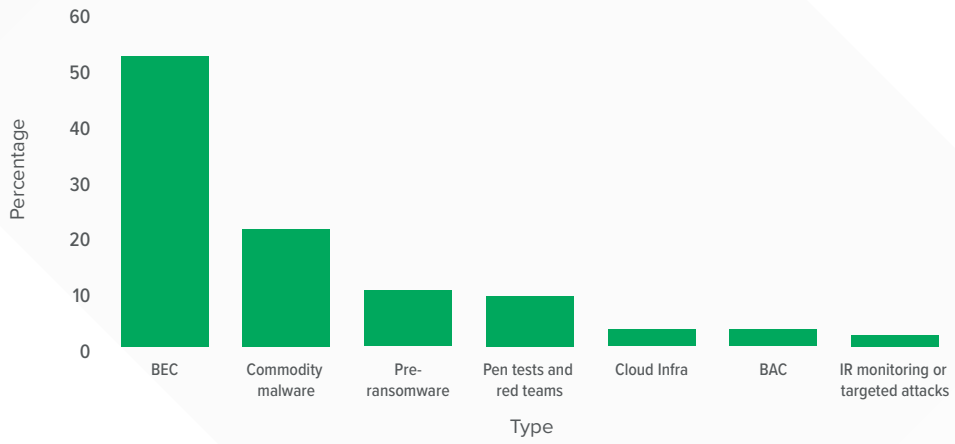

Identity-based attacks accounted for 59% of all incidents handled by our SOC in Q3—up three percentage points compared to Q2.

1 Red team: "A group of people authorized and organized to emulate a potential adversary's attack or exploitation capabilities against an enterprise's security posture." (NIST Computer Security Resource Center)

Purple team: "A purple team is the combination of both offensive and defensive cybersecurity professionals, who perform their responsibilities as a single unit. The security departments of most organizations are made up of a red and blue team." (MakeUseOf.com)

Incident leads by tech type

Compared to Q2, we saw more incidents beginning with identity-based alerts and fewer incidents coming from EDR. In Q3, 59% of all incidents our SOC identified began with an initial lead from an integration with a cloud infrastructure or identity provider, an uptick of five percentage points over Q2. On the other hand, 33% of incidents started with an initial lead from an EDR integration, representing a five percentage point drop from Q2. While network detection and response (NDR) and security information and event management (SIEM) account for only 9% of initial leads into Q2 incidents, these technologies provide SOC analysts with significant investigative capabilities and power orchestration in the Expel Workbench™.

Chart 2: Initial leads into incidents by tech integration type

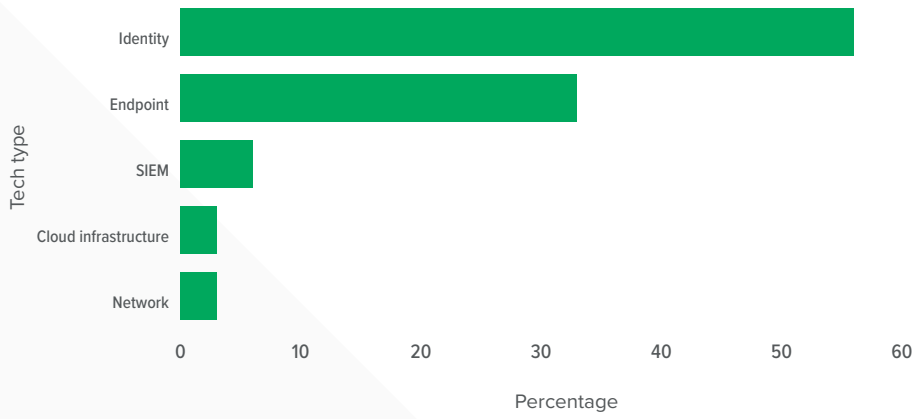

Alert and investigative orchestration

To improve our SOC's scale and quality, we automate a lot of our analysts' repetitive tasks—things like "bounce this file hash against VT," or "tell us the user's role in the organization." This frees analysts up to focus on risk-based decisions for our customers vs. spending time fighting with a query language to retrieve results.

How does orchestrated automation contribute to freeing up analysts? Every alert that became an incident included key investigative data gathered through automation *before* it was sent to our SOC for review. When analysts spend less time on manual tasks, it improves scale and levels up quality by standardizing investigative steps.

Response orchestration

Orchestration not only improves scale and quality, it also accelerates remediation. When our SOC identifies an incident, analysts investigate to uncover the scope and create remediation actions to reduce risk. Expel Workbench can automatically complete remediation actions for our customers, such as containing a host, disabling an account, removing phishing emails, or adding attacker IOCs/hashes to a "deny" list.

In Q3, the median time to complete a remediation action not automated through orchestration was 99 minutes. What happens when a remediation action is automated via orchestration? That median time drops to 15 seconds—a 99.7% improvement.

What happens when a remediation action is automated via orchestration? That **median time drops to 15 seconds—a 99.7% improvement**.

Response speed

The median alert-to-fix time for critical incidents our SOC handled in Q3 was 26 minutes. (Alert-to-fix is the total time from when an alert landed in Expel Workbench to when we tell our customer what they need to do or take that action on their behalf.)

Incident volume by day of the week

In our Q2 report, we found that Tuesday and Friday were the busiest days of the week, with a relatively low incident count on the weekends. This quarter, it's still relatively slow on the weekends, with a more gradual upward trend and peak incident activity on Tuesday and Wednesday (which together account for 41% of all incidents). The weekends combined for 8% of our incidents last quarter, but this quarter weekend incidents rose 75 percentage points to account for 14% of the total. It looks like cyber criminals put in a few extra overtime hours more often.

The key takeaway? Our data suggests we can continue to expect a relatively low volume of incident activity on Saturdays and Sundays compared to typical business working days. However, organizations should consider seasonality in security operations work and lower (but not zero) work volume on weekends. This helps staff appropriately and ensure the right escalation procedures are in place (if needed)—particularly as we head into the busy holiday season.

Chart 3: Incident activity by day of the week

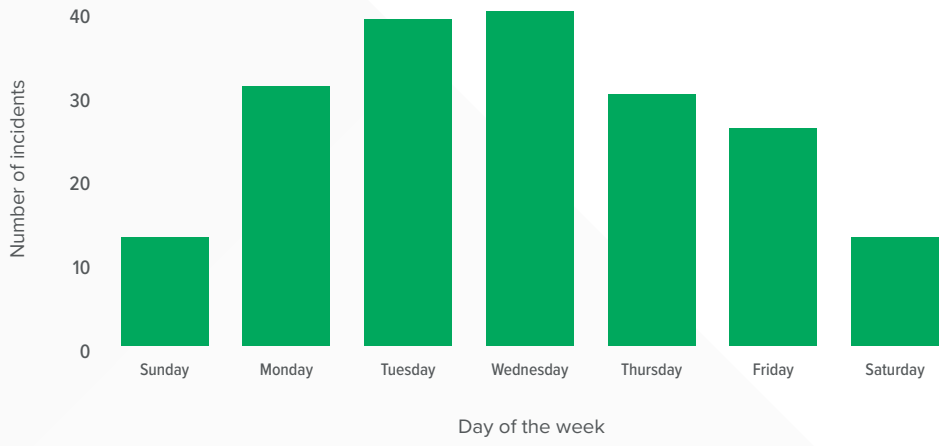

Organizations should **consider seasonality in security operations work** and lower (but not zero) work volume on weekends.

How do we **use automation**?

Learn more

Business email compromise (BEC)

TL;DR

BEC accounted for 52% of all incidents, with 100% occurring in Microsoft 365 (formerly Office 365) for the second quarter in a row. Fourteen percent bypassed MFA using legacy protocols, a decrease of five percentage points compared to Q2. This is almost certainly related to attackers preparing for Microsoft's long-awaited disabling of Basic Auth for Exchange Online, which went into effect on October 1, 2022.

Over half of the incidents we respond to are BEC, so it's safe to say BEC is once again the top threat to our customers. All Q3 BEC attempts occurred in Microsoft 365—we didn't identify any incidents in Google Workspaces. We believe that's due to Google Workspaces having more stringent security settings configured by default. We're watching closely to see if that changes with the recent Basic Auth change for Microsoft 365.

For context, we monitor roughly twice the number of tenants for Microsoft 365 as we do Google Workspaces. But the fact that we didn't identify a single BEC attempt in Google Workspaces yet again is certainly...interesting. Is this the "Mac malware doesn't exist" of webmail? At the minimum, it's an intriguing mystery.

Heads up: If you're in the U.S. and think you only need to closely monitor for non U.S.-based IPs attempting to log in to your environment, know that almost half of the BEC attempts and successful BEC compromises we see originate from U.S.-based IP addresses. To be exact, 48% of the BEC attempts and 46% of the times an attacker successfully accessed an account in Q3 originated from an IP address with a U.S. geolocation.

Also, 100% of authentication attempts originating from the U.S. came from an IP associated with a VPN or hosting provider. This tactic offers a number of advantages for remote attackers. They're more likely to bypass conditional access policies for source countries that either force the user into an MFA challenge or even flat out block the login. If attackers gain access to the account by harvesting user credentials instead of brute force or another method, they can also harvest the user's IP (and therefore geolocation). For authentications, it's vital to alert based on the IP organization as well as VPN enrichment services. We recommend ipinfo.io and spur.us.

For authentications, **it's vital to alert based on the IP organization** as well as VPN enrichment services.

Chart 4: Percent of BEC attempts by country

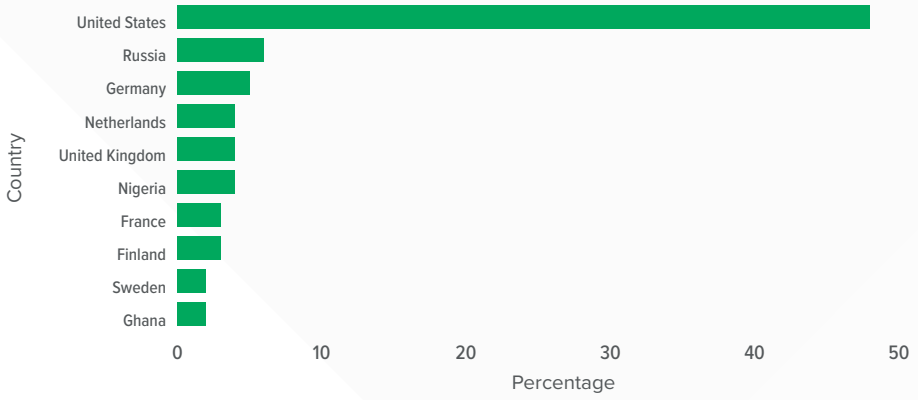

Chart 5: Percent of successful BEC compromises by country

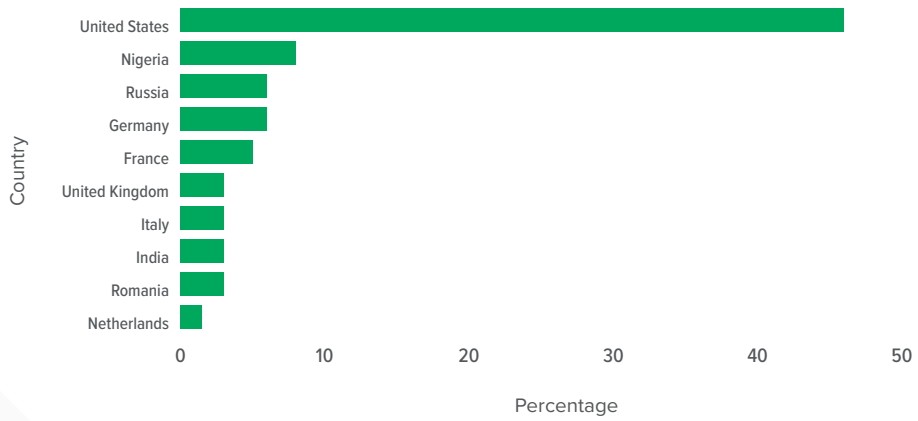

The key takeaway? Threat actors use a wide range of network infrastructure to bypass conditional access policies, including IP addresses associated with VPN services and hosting providers in the same geolocation as the compromised user. Single-factor authentication backed by conditional access policies simply isn't enough to prevent unauthorized access. Our recommendation? Combine conditional access policies with MFA in Microsoft 365. We strongly recommend phish-resistant Fast ID Online (FIDO) security keys.

BEC attempt trends

In Q3, about 14% of BEC incidents involved an attacker attempting to bypass MFA through basic authentication protocols. This is down five percentage points from Q2.

Last quarter we talked about BEC attempts leading to payroll fraud. That's still a thing, and we just finished an integration with Workday to spot this activity at the source. As a refresher, more attackers seem to be skipping the hassle of tricking payroll and going for a more distributed model of changing employees' direct deposit information in a payroll system. The attacker then accesses the compromised user's email account and creates rules to hide the activity.

BEC targeting by industry

Chart 6 shows the percentage of BEC attempts our SOC identified in Q3 across the industries within our customer base. The data shows that BEC fraud isn't an industry-specific problem. Last quarter, criminals targeted retail and transportation most often. Those two industries remain popular targets, but this quarter, government organizations rose to the top—coming in at 31%. That's more than twice as many BEC attempts as the entertainment industry, which came in at 14%. 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 anyone.

Single-factor authentication backed by conditional access policies **simply isn't enough to prevent unauthorized access**.

Chart 6: BEC targeting by industry

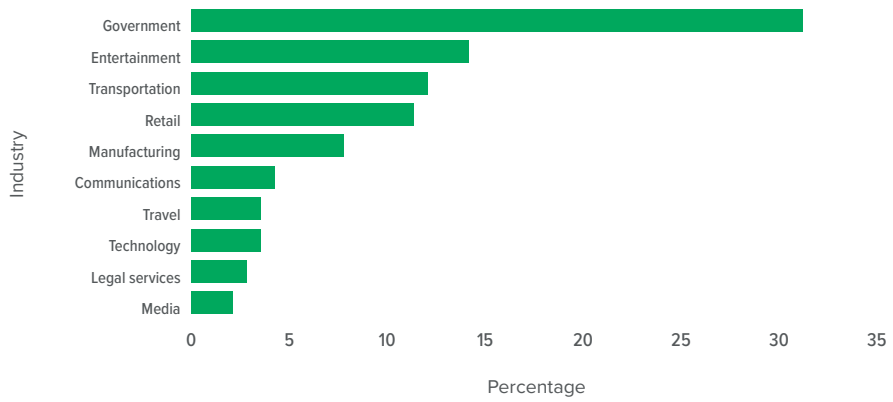

BEC targeting by role

In Q3, threat actors targeted individual contributors in 60% of BEC attempts. Those in management roles accounted for 29% and executives made up 11%.

The bar graph in Chart 7 below shows compromise 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 and at the executive level.

Even though executive-level targeting accounted for only 11% of BEC attempts in Q3, it's worth noting that it had the highest success rate. This may be related to the comparatively low rate of mitigation, like MFA and conditional access policies configured for executive accounts, relative to management and individual contributor accounts.

Chart 7: BEC targeting by role

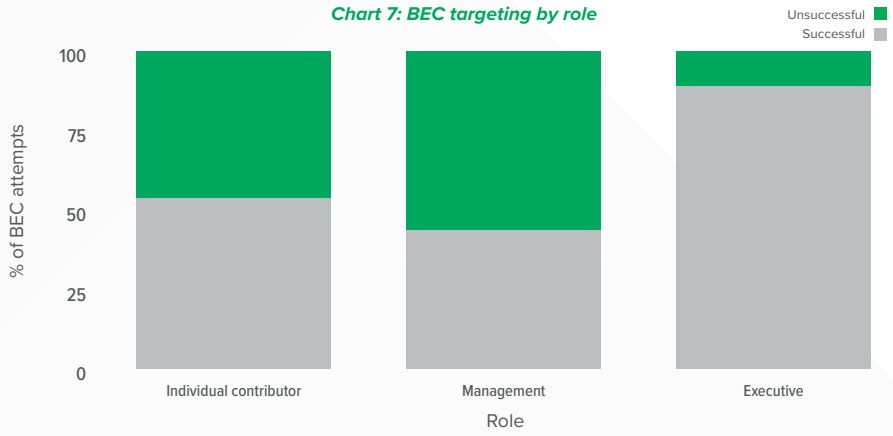

BEC targeting frequency

About 25% of our customers experienced at least one Microsoft 365 BEC attempt in Q3.

With respect to BEC targeting frequency (how often threat actors target an organization), 36% of targeted customers saw more than three attempts. For one government organization, our SOC identified 44 BEC attempts in its Microsoft 365 tenant alone—that's roughly once every other day. The takeaway? BEC remains a constant threat to organizations—regardless of industry or role.

HOW TO PROTECT YOUR ORGANIZATION

1. Make sure you're running MFA wherever possible, using phish-resistant FIDO security keys.
2. Disable legacy protocols like IMAP and POP3. They don't support any sort of modern authentication, which means an attacker can bypass MFA completely by using an IMAP/POP3 client. Once you turn these off, strongly consider disabling Basic Authentication to prevent any pre-auth headaches for your Microsoft 365 tenants.
3. Next, implement extra layers of conditional access for your riskier user base (such as executives or employees with access to sensitive data) and high-risk applications. You can create a conditional access policy to require MFA registration from a location marked as a trusted network, preventing an attacker from registering MFA from an untrusted network.

Want to learn more about **how Expel can stop BEC**?

Let's chat

Business application compromise (BAC)

TL;DR

Just like last quarter, BAC attempts led to about 3% of incidents we saw; these occurred in Okta, Ping Identity, and OneLogin. The good news: only about half the incidents resulted in the attacker successfully gaining access to the account, with the other half being stopped by MFA or conditional access policies. The bad news: MFA and conditional access were configured for over 80% of the successful compromises, but the attacker successfully tricked the legitimate user to satisfy the MFA request. This is way up from last quarter, when only 14% of successful compromises came from repeated push notifications. The takeaway? Similar to phishing training, end-user training on the importance of reporting suspicious MFA activity is a must.

MFA push notification fatigue attacks are increasing because they're working. As more are turning to cloud access identity providers like Okta, Ping Identity, or OneLogin to provide a more convenient single sign-on (SSO) experience for their employees, companies must be aware of the potential risk. With SSO, attackers only need to steal one credential for access to more than just a user's email.

Thirty percent of the incidents we detected involved threat actors using a VPN infrastructure located in the U.S., likely to bypass conditional access policies implemented by organizations. Additionally, in 30% of the BAC attempts we detected, the attacker either used a mobile device or spoofed the user-agent string to suggest it was a mobile device—likely to blend in with the user's normal authentication behavior.

BAC source IP addresses by country

Thirty percent of the BAC incidents our SOC identified involved threat actors using an IP address that was located in the U.S. VPN providers hosted 100% of these IP addresses. Threat actors were likely trying to bypass conditional access policies that restrict access from certain countries and blend in with normal user behavior.

Chart 8: Locations of IP addresses used in BAC incidents

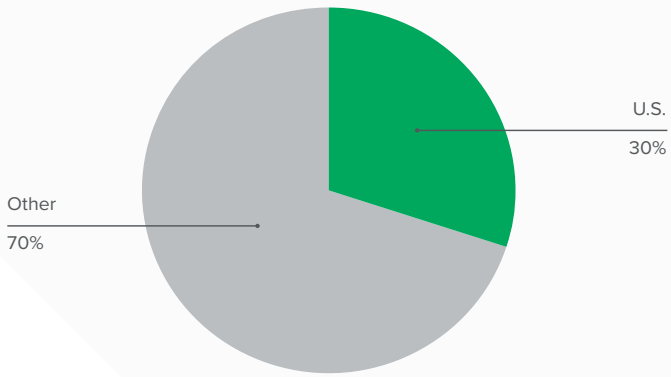

Thirty percent of the BAC incidents our SOC identified **involved threat actors using an IP address that was located in the U.S.**

BAC source IP address by type

The pie chart below (Chart 9) shows the distribution of IP address infrastructure types threat actors used. They employed Internet Service Providers (ISPs) 40% of the time, and an even split between VPN and hosting providers—both 30% of the time.

Chart 9: BAC source IP address by type

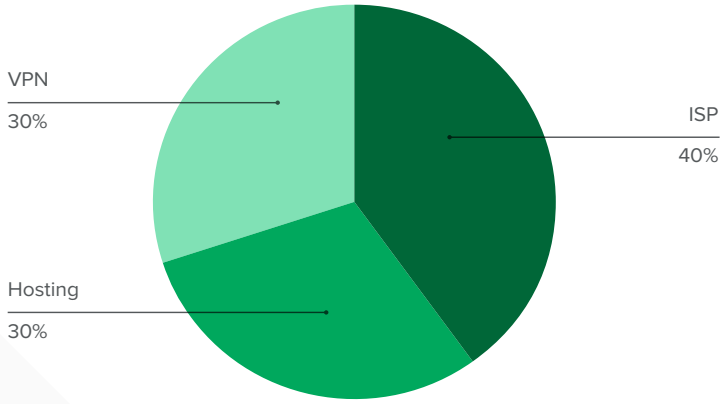

How to spot it

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

- Alert on successful SSO or SaaS authentications where the source IP address is a proxy, the device is new, or when the source of activity is a known data center, hosting provider, or VPN service.
- Alert on authentications that represent geo-impossible travel using the Velocity policy. For Okta specifically, be sure to configure the Okta Velocity policy with a speed that represents geo-impossible travel (e.g., 700 km/h for a commercial airline).
- Alert for multiple Okta sessions from the same user with multiple, non-mobile operating systems.
- Alert for potential brute force Duo push requests.
- Alert on Duo authentications where the access and authentication IP addresses represent a distance that's likely geo-infeasible (that is, they're really far apart).
- Alert when Duo blocks an anomalous push notification, as this can indicate a compromise of a username and password combination.

HOW TO PROTECT YOUR ORGANIZATION

1. To stop MFA push notification fatigue attacks, organizations can disable push notifications in favor of a PIN or a Fast Identity Online (FIDO) compliant solution.
2. If that's unrealistic, control push notifications using number matching—a setting that requires the user to enter numbers from the identity platform into their MFA app to approve the authentication request.
   - Microsoft, Duo, and Okta all support this feature.
3. Implement a pre-auth policy for network zones in Okta.
4. Consider blocking access to Okta from suspicious network zones based on IP address(es), autonomous system numbers (ASN), IP type, or geolocation.
   - Why? Clients from blocked zones can't access any Okta URLs and requests are automatically blocked before authentication.
5. Deploy Okta's adaptive multi-factor authentication (AMFA).
   - AMFA 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 spot identity threats**?

Let's connect

Pre-ransomware

TL;DR

Pre-ransomware incidents accounted for 10% of all incidents we detected—consistent with Q2. Threat actors continue to use zipped Javascript and ISO files to deliver malware for initial access, accounting for 70% of all pre-ransomware incidents. In addition, threat actors delivered 10% of pre-ransomware malware through removable media devices. We also continued to see a decline in the use of Microsoft Excel files, from 9% in Q2 to 8% in Q3.

Our SOC attributed 10% of incidents to pre-ransomware activity. If we hadn't detected and remediated this activity, the threat actor would likely have ransomed the target organization.

The data focuses on the deployment of malware we've linked to potential ransomware operations. This includes initial droppers/downloaders and backdoors enabling remote access that threat actors might sell to a ransomware affiliate.

Initial entry

The top attack vectors used by ransomware groups to gain initial entry were:

1. Zipped JavaScript files (46% all pre-ransomware incidents)
2. Zipped ISO files (26%)
3. Removable media (10%)
4. Excel 4.0 macros (8%)

The bar chart (Chart 10) below shows the pre-ransomware initial attack vector for Q3 2022. The height of the bar represents, as a percentage, how often a particular attack vector led to pre-ransomware incidents.

Chart 10: Pre-ransomware malware attack vectors

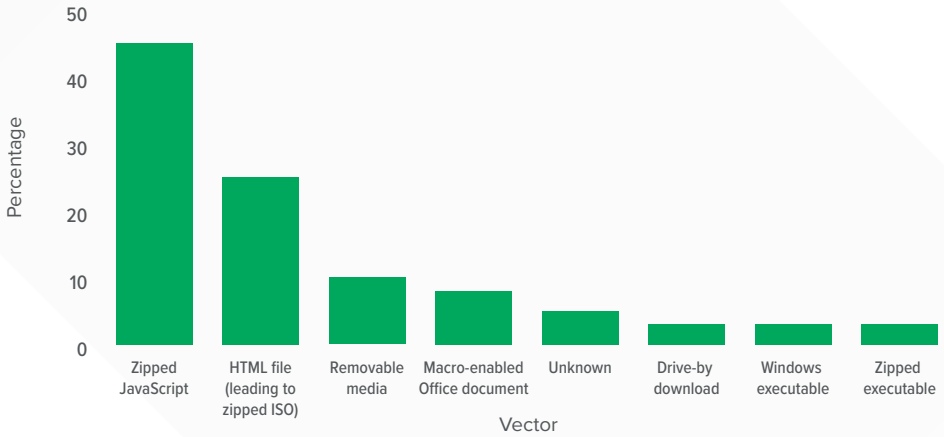

Our SOC continued to see the Q2 trend of threat actors using zipped JavaScript and ISO files to deliver malware to gain initial access. Way back in Q1 (when Microsoft announced its plans to disable Excel 4.0 macros by default in Q3),

Threat actors continue to use **zipped Javascript and ISO files to deliver malware for initial access**, accounting for 70% of all pre-ransomware incidents.

a macro-enabled Microsoft Word document (VBA macro) or Excel 4.0 macro was the initial attack vector in 55% of all pre-ransomware incidents. In Q3, we saw Microsoft Excel 4.0 documents used to deliver malware just 8% of the time.

A notable change in Q3 was the increase in removable media as an attack vector. In Q2, about 4% of the incidents associated with pre-ransomware stemmed from removable media, but that jumped to 10% in Q3. We attribute this rise to the spread and delivery of the **RASPBERRY ROBIN** family of malware, first identified in late Q2, that's been linked to pre-ransomware operations.

Ransomware targeting by industry

Ransomware threat groups targeted organizations in manufacturing the most, accounting for 18% of all the pre-ransomware incidents we detected. Legal services and insurance both accounted for 13% of incidents, and these three verticals made up a total of 44% of all ransomware incidents we detected in our SOC.

Chart 11: Ransomware industry targeting

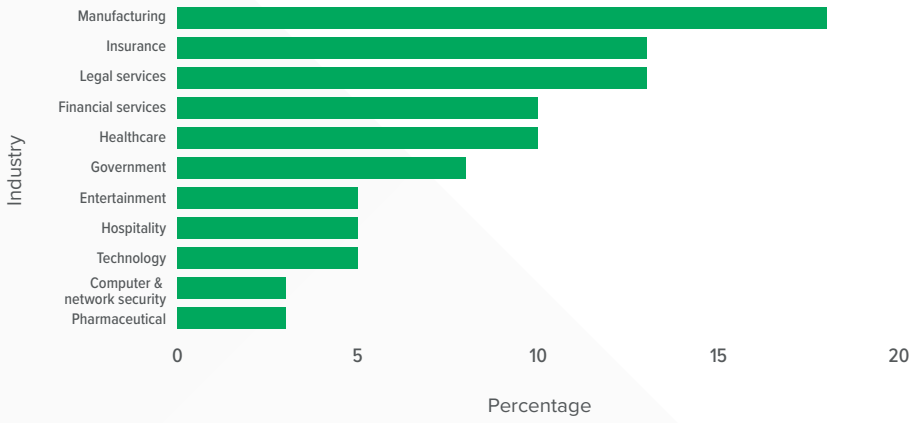

Initial access actions

Here are a few examples of native Windows OS binaries used for initial access:

- Windows Calculator (calc.exe) launches the register server (regsvr32.exe) to load a malicious dynamic link library (DLL) file.
- The Microsoft Excel process (excel.exe) launches the register server (regsvr32.exe) to execute a file from the active Windows user profile.
- A scripting process other than PowerShell (like wscript.exe) launches a PowerShell process with encoded commands.
- 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.
  - Launches Windows command prompt (cmd.exe).

HOW TO PROTECT YOUR ORGANIZATION

1. Disable auto-play functionality for all removable media devices.
2. Unregister ISO file extensions in Microsoft Windows Explorer.
   - Windows will no longer recognize ISO files and double-clicking won't result in program execution.
3. 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.
4. Disable Excel 4.0 macros. Microsoft has disabled Excel 4.0 macros from the internet by default, but it's important to understand if they're still enabled for your organization.
5. IT administrators should set policies that block active content in Office docs that arrive by email. Microsoft also started providing more granular controls for macros, ActiveX content, and Office add-ins in emailed Office docs as of February 2022.
6. Don't expose remote desktop protocol (or any other service you don't need to) directly to the Internet.

Phishing

TL;DR

Sixteen percent of phishing emails submitted to our team for review were malicious, and 6% were simulations run by the customer organization. The top subject line themes for malicious emails were no subject line, followed by "Invoice," "Order confirmation," "Payment," and "Request." Our data shows that social engineering themes that create urgency, fear of missing out (FOMO), or potential financial loss, spur the most action (open, click, interact).

Customers send our phishing team suspicious emails to determine if they're malicious or just unwanted spam. This gives us unique visibility into the various phishing attacks launched to perpetuate BEC fraud, gain initial access to a target organization, or even phish for AWS account root user credentials.

Sixteen percent of the phishing emails our analysts reviewed were malicious. Six percent were phishing simulations run by the customer organization. These malicious emails contained links to download malware, links to credential harvesting sites, or attachments that contained malware droppers.

Top malicious subject line themes

While the text may change, our data shows that threat actors love a good theme when it comes to subject lines. We've taken the most common ones submitted to our SOC and distilled them down into themes. The most common offender? No subject line. The rest are the themes you'd expect—invoice, order, payment, urgent, etc.

Table 1: Top ten subject line themes in malicious emails

Top subject lines Note: These aren't the actual subject lines, but the originals did include some component of the theme.	**Percentage**
<blank>	57%
Invoice	9%
Order confirmation	9%
Payment	8%
Request	7%
Document	5%
Urgent	2%
Deposit	2%
Payroll	1%
Shared file	1%

Our data shows that **social engineering themes that create urgency, fear of missing out (FOMO), or potential financial loss, spur the most action** (open, click, interact).

Submitter accuracy

We investigated submissions from tens of thousands of unique submitters from our customers over the past quarter. Of those submitters, 69% strictly submitted benign emails. Obviously there is room to hone this enthusiasm, and if we remove those users and focus strictly on the 31% of users who submitted at least one malicious email, we find something interesting: they're good at spotting phishes.

Sixty-six percent of users who submit at least one malicious email correctly have perfect accuracy.

Chart 12: Submitter accuracy for emails submitted to our SOC

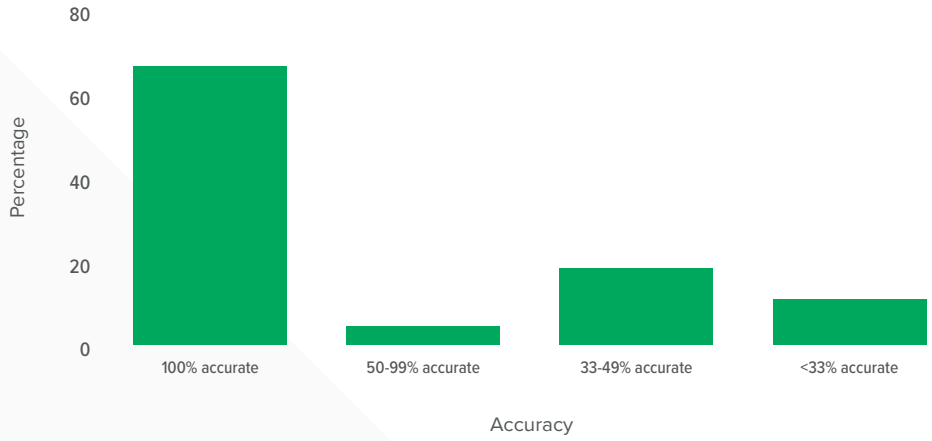

The key takeaway? While most submitters have never submitted a malicious email, the second largest group *only* submits malicious emails.

HOW TO PROTECT YOUR ORGANIZATION

1. Make sure you're running MFA wherever possible, 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.

Penetration testing, red teams, and purple teams

TL;DR

Nine percent of incidents were authorized penetration tests, red team, and purple team exercises. Red and purple teams preferred Cobalt Strike for post-exploitation (post-ex) and command-and-control (C2); they used a combo of RDP, WMI, and PsExec to move laterally. Red teams most often choose ScoutSuite for cloud security assessments in AWS.

Nine percent of incidents detected by our SOC were authorized penetration tests, red team, and purple team exercises, which allow organizations to test their security controls, remediation processes, and investigative capabilities.

One of the notable findings in our data for Q2 was that 22% of red team engagements were performed in AWS environments. In Q3, that number fell to 12%.

Effective red and purple team exercises emphasize detection, response, *and* remediation.

Red team and purple team engagements provide a good reality check of an organization's detection and response capabilities. They help determine a few things about how you'd fare against a threat. For starters, could your technology detect it? Could your SOC identify the initial entry point? And is there anywhere else a red team used a compromised account?

We averaged 23 alerts per red team, and our SOC performed an average of 16 investigative actions when chasing down a red or purple team. Alerts provide leads, and investigation uncovers the scope. If you can get an initial lead and start investigating early enough, you can make it really hard for a red team to succeed.

These engagements also stress test your remediation processes. On average, our SOC recommended six remediation actions per red team. This included containment for batches of hosts, password resets for compromised accounts, and blocking file execution of known payloads. Why does this matter? Practicing the detection and investigation is great, but the whole point is to *stop* the attacker. When the real thing happens, you must react immediately. (Saturday at 3 a.m. is *not* the time to find out you don't know the phone number of the security team member who blocks domains...) Actually walking through a simulated attack with your team always seems to shake out a potential show-stopping issue or two.

Some general themes in our data:

- Red and purple teams preferred Cobalt Strike as the post-ex and C2 framework.
- Multiple red teams used CrackMapExec for enumeration and lateral movement via Server Message Block (SMB).
- Red teams used LaZagne and Mimikatz to steal Windows credentials.
- Red teams moved laterally via remote desktop protocol, Windows Management Instrumentation (WMI), and through ImPackets' PsExec module.
- On the cloud infrastructure side, red teams preferred ScoutSuite to perform assessments in AWS.

HOW TO ENABLE YOUR ORGANIZATION

1. Red team exercises should emphasize response. Talk about remediation ahead of time. Ask hard questions like, "What would we do if that account was compromised?"
2. Review your incident response (IR) plan with the team. It's important to build muscle memory around your IR process *before* a red team exercise.
3. Use an MSSP or MDR? Chat with them. Understand the rules of the road for responding to red team activity. One of your red team goals likely includes assessing your MSSP/MDR. That's great, but understand what you can expect before you get started.

Want to learn more about **how Expel responds to red and purple team exercises**?

Let's chat

Incident response (IR) monitoring

TL;DR

We partner with IR consulting firms to provide 24x7 SOC services during IR engagements. The Expel Workbench helps IR consultants get around-the-clock monitoring up and running quickly for new engagements. Technology onboarding requires a few simple steps and Workbench provides a seamless experience so our SOC has the situational awareness needed to be effective.

A big part of our value-add is our SOC's ability to triage alerts, investigate, and complete remediation actions on behalf of our IR partners. This creates space for them to focus on the overall investigation and not get distracted chasing down activity unrelated to the breach.

Chart 13 illustrates the rough percentage of the different kinds of IR engagements we supported in Q3. Ransomware still tops the list, accounting for about 75% of the total. We recently started supporting BEC engagements, and it roughly matched the APT engagements we supported in Q3, with both representing about 12.5%.

Chart 13: IR engagement types

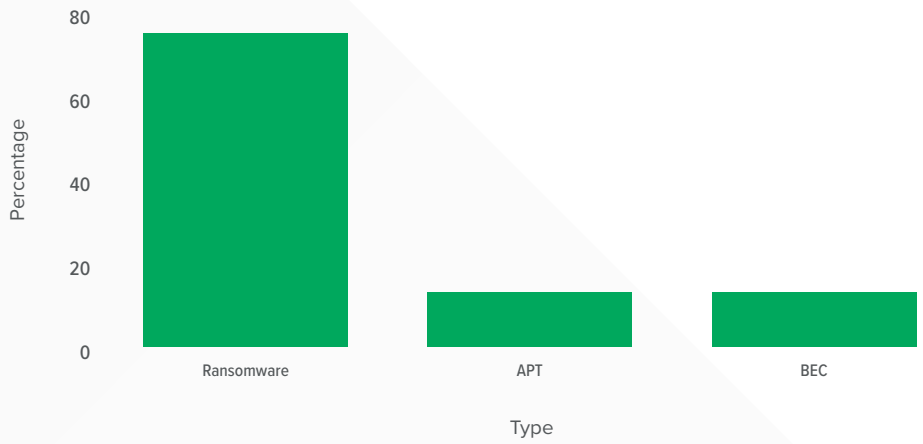

Our SOC tends to spend most of its time chasing attackers in the earlier phases of the attack lifecycle (a good thing). Partnering with IR consulting firms gives us visibility into what attackers are doing in later phases, providing new experiences for our analysts and leading to improved detection and strategic recommendations for our customers.

HOW TO STRENGTHEN YOUR IR CAPABILITIES

1. Consider an IR retainer to reduce your incident response time. You can operate under pre-negotiated terms and communication channels to get a response started quickly.
2. Review your IR plan with the team. It's important to build muscle memory around your IR process *before* an incident.

Want to **partner with Expel for your IR engagements**?

Let's connect

Looking ahead to Q4

Microsoft's announcement that it would block macros by default in Office applications changed hacker behavior on the endpoint. Attackers were forced to abandon tried-and-true methods and adapt. Will Microsoft disabling Basic Auth for Exchange Online have a similar effect? Almost certainly. In fact, we've likely seen some of those impacts already. These include increased VPN usage for geolocated IPs to blend in with the target account and bypass conditional access rules, an increased reliance on OAuth attacks, and no doubt some new tricks, as well. And if that fails, there's always the numbers game of spamming users with MFA pushes until someone approves.

Speaking of which, we think it will become increasingly important to configure MFA in a way that prevents an attacker from spamming users with push notifications and to train users to identify and report suspicious activity. By the time the attacker has gotten to an MFA prompt, the user's credentials are often already compromised. Even if the attacker doesn't manage to get in right away, there's a chance to find another avenue with the user's credentials. Users need to know the tools available to them to report suspicious MFA requests to their security teams. We also anticipate adoption of MFA number matching features to reduce risks associated with MFA fatigue attacks. When that happens, we expect to see a shift in MFA bypass techniques. But until then, this technique will continue to be used more and more in identify-based attacks.

We thought BEC would remain the top threat organizations face in Q3, and we anticipate that it will be the same for Q4. However, we'll continue to see more BEC attempts as a means to access payroll management systems vs. a path to perpetrate wire transfer fraud. We also suspect that we will slowly but surely see attackers put more effort into compromising Google Workspace accounts, as Microsoft 365 makes changes to increase the security of its default settings.

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. Expel offers managed detection and response (MDR), remediation, phishing, and threat hunting. For more information, visit our website, check out our blog, or follow us on LinkedIn and Twitter.

WANT TO LEARN MORE?

- Learn about the problems we solve
- Watch a video demo
- Subscribe to our blog
- We're hiring! Find the right role for you
- See what Expletives say about working at Expel

Authors: Ben Brigida, Jonathan Hencinski, David Blanton, and Brian Badorrek
Contributors: Joshua Chou, Elisabeth Weber, Ray Pugh, and Myles Satterfield