ARCTIC WOLF LABS

Principle networks

ARCTIC WOLF LABS

THREAT REPORT 2024

TABLE OF CONTENTS

EXECUTIVE SUMMARY

This edition of the Arctic Wolf® Labs Threat Report draws upon the first-hand experience of Arctic Wolf's security experts, augmented by research into the cybercrime ecosystem and additional credited data sources.

Here are a handful of the top takeaways:

Ransom demands surge 20% year-over-year to $600,000 USD

Continuing a multiyear trend, the median initial ransom demand in incidents we investigated rose to $600,000: a year-over-year increase of 20%.

And there are worrying signs that 2024 will be especially volatile, as ransomware groups expand their list of targets, and explore new pressure tactics in response to increasingly effective international law enforcement efforts and the growing momentum of refuse-to-pay initiatives.

Ransomware is 15x more likely than business email compromise to lead to an incident response engagement

Ransomware attacks are feared by organizations large and small, and with good reason — the damage and disruption they cause is responsible for immense losses above and beyond the ransom itself.

Attempts to recover these losses through cyber insurance often lead to formal incident response (IR) engagements, as insurers seek to understand the details of an attack.

Despite business email compromise (BEC) incidents outnumbering ransomware incidents by a factor of 10 (as reported by the FBI), nearly half (48.6%) of IR engagements conducted by Arctic Wolf are in response to ransomware. In fact, combining the FBI's figures with our own suggests that a ransomware incident is 15 times more likely than a BEC incident to lead to an IR engagement.

Business email compromise is the unacknowledged, unyielding threat to global organizations

Ransomware garners more headlines, but BEC incidents are effective and much easier to execute. Plus, only the most severe BEC incidents — for instance, those with account compromise or other intrusion actions — typically lead to a full IR engagement.

Nevertheless, BEC incidents accounted for 29.7% of the total incidents investigated by Arctic Wolf® Incident Response during this reporting period, underscoring how much of an everyday threat they remain for today's organizations.

Nearly half of all successful attacks are driven by credential reuse

In recent years, threat actors have become increasingly adept at acquiring and using credentials.

Most BEC incidents — whether involving account compromise or limited to spoofing or masquerading — can be traced to phishing, while 46.3% of non-BEC attacks are driven by credential reuse. More specifically, 39% of non-BEC incidents Arctic Wolf investigated involved an attacker using credentials to log into an external remote access application, while another 7.3% of non-BEC incidents leveraged previously compromised credentials to gain direct access to a victim's environment via other asset types.

Organizations can strengthen their security posture by enforcing robust identity controls, including strong multi-factor authentication (MFA) and passwordless authentication techniques, and by implementing modern identity and access management (IAM) infrastructures.

Threat actors succeeding en-masse by exploiting 2-year-old vulnerabilities

In 29% of non-BEC incidents Arctic Wolf investigated, the attacked exploited a vulnerability. Notably:

Nearly 60% of these incidents exploited a vulnerability identified in 2022 or earlier, meaning organizations had anywhere from months to years to patch the affected system or remove (or further safeguard) its external access.

Only 11.7% of these non-BEC incidents — or 3.4% of incidents, overall — featured a zero-day exploit.

Patching is clearly an effective way to prevent incidents, and while it can seem like an overwhelming activity, a little prioritization can go a long way.

More than half of the incidents Arctic Wolf investigated involved at least one of 10 specific vulnerabilities — taking these 10 off the table will make life harder for threat actors.

Similarly, although there's a long list of tactics, techniques, and procedures (TTPs) available to threat actors, a relatively small number show up repeatedly in Arctic Wolf's engagements.

INTRODUCTION

Most cyber attacks ultimately fail to achieve their goals — that is, for every incident that progresses to the actions on objectives stage, many others are stopped by layers of cybersecurity technologies, members of the workforce recognizing suspicious activity, or, as a last line of defense, by internal or third-party security personnel responding to alerts.

While there's value in studying how a combination of people, processes, and technology detected and stopped an intrusion earlier in the attack chain, this report directs its attention on cyber attacks that succeed — the ones in which a threat actor accomplishes their goals.

By focusing on successful cyber attacks, we aim to:

01 **Highlight which attack types are most responsible for severe incidents**

02 **Uncover the TTPs that are allowing threat actors to evade detection long enough to damage and disrupt victims**

03 **Raise awareness of the cybersecurity practices that are needed to prevent, detect, and recover from such incidents**

Gaining insights by studying severe incidents

This report focuses primarily on hundreds of digital forensics and incident response (DFIR) engagements conducted by the Arctic Wolf Incident Response team.

The vast majority of these engagements were initiated as part of cyber insurance policies, through our partnerships with insurance providers and privacy law practitioners.

Consequently, these incidents typify cyber attacks that are so severe (i.e., damaging, disruptive) that they led to insurance claims — making them ideal study subjects in our aim to better understand the most dangerous threats.

Ransomware (and data extortion) and business email compromise together account for nearly 80% of the incidents we investigated.

Accordingly, Part One is devoted to examining these particular threats in detail. In Part Two, we explore the root causes behind such incidents before diving more deeply into vulnerabilities and TTPs in Part Three—incorporating data from the Arctic Wolf® Managed Detection and Response (MDR) solution. While we intersperse some security recommendations throughout the report, we reserve most threat mitigation and management guidance until Part Four, before closing off with some final conclusions.

Data sourcing and methodology

Unless otherwise stated, all data and analysis within this report pertains to the period beginning November 1, 2022, and ending October 31, 2023, with information from three main sources:

Arctic Wolf Incident Response (IR)

DFIR engagements performed by the Arctic Wolf Incident Response team. As noted above, these engagements are typically initiated via cyber insurance and privacy law. Cyber insurance is a valuable risk management approach for any organization, however we recognize that certain industries are more likely to have coverage than others, and that our sample cases will reflect this distribution.

Arctic Wolf® Labs

Information and insights provided by Arctic Wolf® Labs, which brings together elite multi-discipline security professionals to deliver cutting-edge threat intelligence and security research, develop advanced threat detection models, and drive continuous improvements in speed, scale, and efficacy.

Arctic Wolf Managed Detection and Response (MDR)

The Arctic Wolf® Security Operations Cloud, which processes more than five trillion security events weekly and powers the Arctic Wolf Managed Detection and Response (MDR) solution. While IR engagements are composed of more serious incidents and occur post hoc, MDR emphasizes detection and containment of incidents before they expand in scope and severity — so data gathered through this channel tends to be biased towards initial access vectors and early-stage intrusion actions.

To enable the holistic analysis within this report, all data is aggregated without any identifying characteristics or attributes.

PART 01

ATTACK TYPES

PART 01: ATTACK TYPES

Key Takeaways

- **Ransomware remains the dominant cause of IR cases:** On average, a ransomware attack is 15 times more likely than a BEC incident to lead to an IR engagement. Despite BEC incidents outnumbering ransomware incidents by a factor of 10, overall, nearly half (48.6%) of IR engagements conducted by Arctic Wolf in 2023 were in response to ransomware.
- **Ransom demands continue to rise:** The median initial ransom demanded across all industries was $600,000 USD, a 20% increase year-over-year, and 10 industries received initial ransom demands that were equal to or higher than last year's median.
- **Ransomware groups are becoming more aggressive:** Faced with international law enforcement operations and a rising refusal to pay, groups are expanding their list of targets while also exploring ways to apply even greater pressure to victims — further underscoring the importance of prevention, detection, and recovery.
- **Business email compromise is a pervasive threat:** Ransomware attacks generate more headlines and are behind a higher share of IR cases, but BEC incidents are nevertheless effective and much easier to execute — making them more of an everyday threat to organizations large and small.

At a high level, our engagements can be divided into four incident types:

- **Ransomware (or data extortion), which accounts for nearly half of all incidents**
- **BEC, representing approximately 30% of incidents**
- **Network intrusion, at roughly 15% of incidents***
- **Other (6.9%), which is a catch-all for other threats including malware, fraud, and disruptions**

For context extending beyond our own direct experience, the FBI's most recent Internet Crime Report[1] (see "resources" at end of report) indicates that, overall, reported BEC incidents outnumber reported ransomware incidents by a factor of 10. Combining the FBI's data with our own observations suggests that a ransomware attack is roughly 15 times more likely than a BEC incident to lead to an IR engagement.

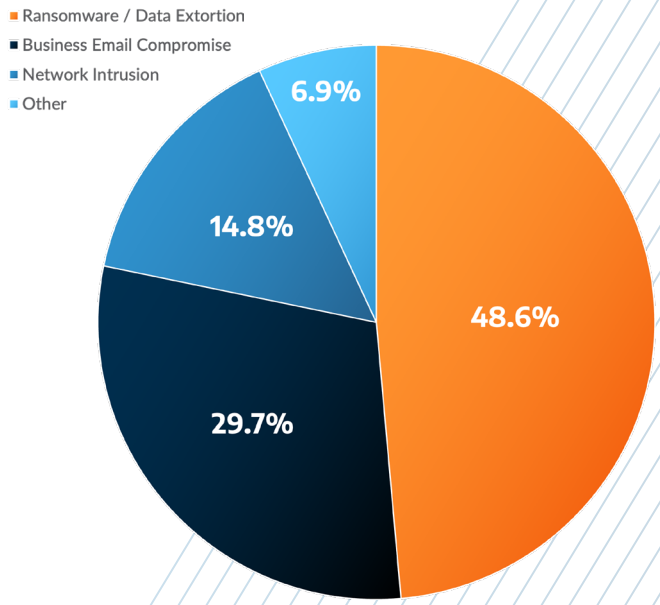

*These are incidents in which some intrusion actions (e.g., lateral movement, privilege escalation) were observed, but the attack was stopped or ended before blossoming into ransomware detonation or BEC.

Ransomware

In recent years, the cybercrime industry has matured and its constituent organizations — including ransomware groups — have grown more sophisticated.

In the ransomware-as-a-service (RaaS) model that has emerged, RaaS operators offer technical resources (e.g., encryption software, leak sites) and branding to independent affiliates who perform the work of compromising and extorting victims — with the proceeds split between affiliates and the operators.

Today, the RaaS ecosystem and affiliate model allows practically any aspiring cybercriminal to participate in attacks, and double-extortion attacks, in which the attacker disrupts operations and threatens to publish exfiltrated data, are the norm. Plus, some ransomware groups and affiliates add additional elements of extortion by directly contacting individuals and organizations with ties to victimized targets.

Meanwhile, remote or hybrid work arrangements are common, extending attack surfaces into home networks, coffee shops, and other locations beyond the control of an organization's IT department. With an uptick in cloud services, more endpoints, unmanaged/BYO devices, and business operations transitioning from analog to digital platforms, stopping ransomware attacks with effective prevention, detection, and response becomes more challenging by the day.

The median ransom demand reaches a new high

Cybercriminals base their initial ransom demand in any particular incident on several factors, including:

- The victim organization's size and financial position, which threat actors use to estimate the organization's ability to pay
- The victim organization's industry, which influences their sensitivity to disruption and negative press
- The scope of the attack, which typically influences the victim's ability to recover and the impact to their operations
- The victim's insurance coverage: Some ransomware groups actively seek out cyber insurance policies in a victim's environment to better inform their ransom demands, typically asking up to the maximum the insurance policy will cover
- The ego and mood of the attacker

Nevertheless, aggregate analysis across hundreds of engagements is still a useful endeavor, as it can reveal trends and shifts in the ransomware economy.

In unwelcome — and unsurprising — news, ransom demands continue to rise:

- The median initial ransom demand associated with incidents investigated by Arctic Wolf Incident Response grew to $600,000 USD — a 20% increase over last year's figure of $500,000 USD*
- 10 industries received initial ransom demands that were equal to or higher than last year's median.

GOING UP...

The median initial ransom demand associated with incidents investigated by Arctic Wolf Incident Response rose 20% year-over-year, to $600,000 USD.

*Median figures are used for comparison purposes because the ransom demand can vary enormously based on the size of the organizations and the scope of the incident. In such a distribution, the median provides the best approximation of what a 'typical' event looks like by limiting the influence of very large and very small outlier incidents.

Median Initial Ransom Demand by Industry

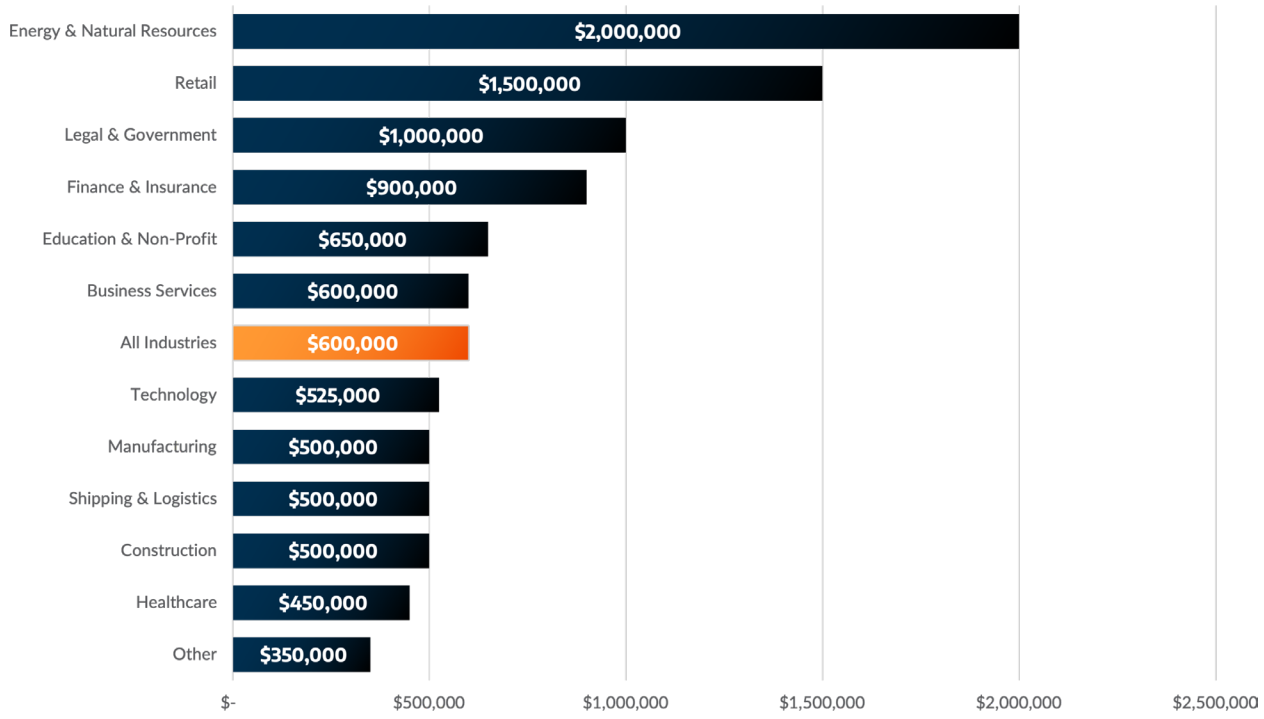

Cybercriminals favor particular industries

We can augment this perspective by examining ransomware group leak sites, but such analysis must take into account that payments have a significant impact on which victims are named on the dark web and which victims remain anonymous.

In particular, ransomware groups use these sites to apply leverage to victims — victims that quickly negotiated with the threat actor and paid a ransom may never be listed on a leak site. Consequently, leak sites have an inherent bias in that they skew towards victims that refuse to pay or are perceived by threat actors as stalling.

For example, manufacturing organizations have more representation on leak sites than any other industry. Threat actors target manufacturers aggressively – they recognize that manufacturers have little tolerance for production downtime – however many manufacturers can maintain production even if incidental or ancillary business systems are temporarily unavailable. In such a scenario, the victim may refuse to pay — at least until they can assess their ability to recover and restore full operations — and their name will undoubtedly appear on a leak site.

In contrast, healthcare organizations are under regulatory pressure to protect the sensitive data they handle, so they may be more inclined to pay a ransom. Similar pressures apply to legal and government organizations. As a result, organizations within these industries are more likely to have robust cyber insurance and to be represented in Arctic Wolf's severe casework, and therefore relatively less likely to appear on dark web leak sites.

MOST-REPRESENTED INDUSTRIES

The industries with the most representation in Arctic Wolf Incident Response ransomware engagements are:

1. **Healthcare**
2. **Education and Non-Profit**
3. **Manufacturing**
4. **Construction**
5. **Legal and Government**

Top 10 Industries Appearing in Leak Sites

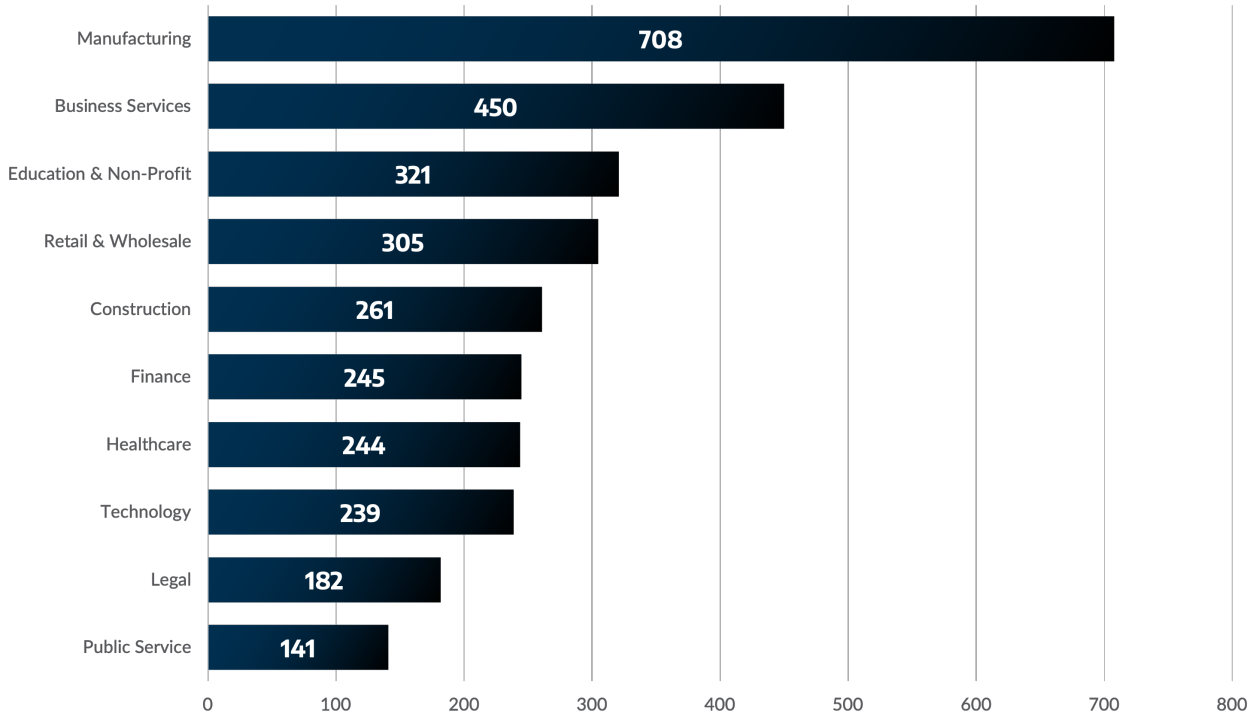

A handful of ransomware variants dominate

In many incidents, investigators can determine with high confidence the ransomware group, or at least the malware variant, behind the attack. Various data points inform such attribution, including:

- Malware samples
- Infrastructure and indicator of compromise (IOC) overlap or reuse
- Post-encryption file extensions
- Ransom message and leak site postings

Focusing on engagements in which the Arctic Wolf Incident Response team confidently attributed an attack to a particular ransomware variant, the five variants we encountered the most were:

1. BlackCat (AlphVM or AlphV)
2. LockBit 3.0
3. Akira
4. Royal
5. BlackBasta

Comparing our variant attribution to groups' leak site data reveals Akira as the greatest outlier. The Akira group gained early attention for their unique leak site[2] and, despite only arriving on the scene in March 2023, has rapidly established itself as a heavyweight in the world of ransomware.

But this year really belonged to LockBit, at least when measured in terms of the sheer volume of incidents they claimed, which was more than double that of BlackCat.

GROUPS, VARIANTS, AND BLURRED LINES

Behind the scenes, the ransomware ecosystem has blurred lines:

- **Individual ransomware groups often work with many different affiliates**
- **Affiliates may use several different ransomware variants — from different groups — concurrently**

10 Ransomware Groups by Claimed Victims

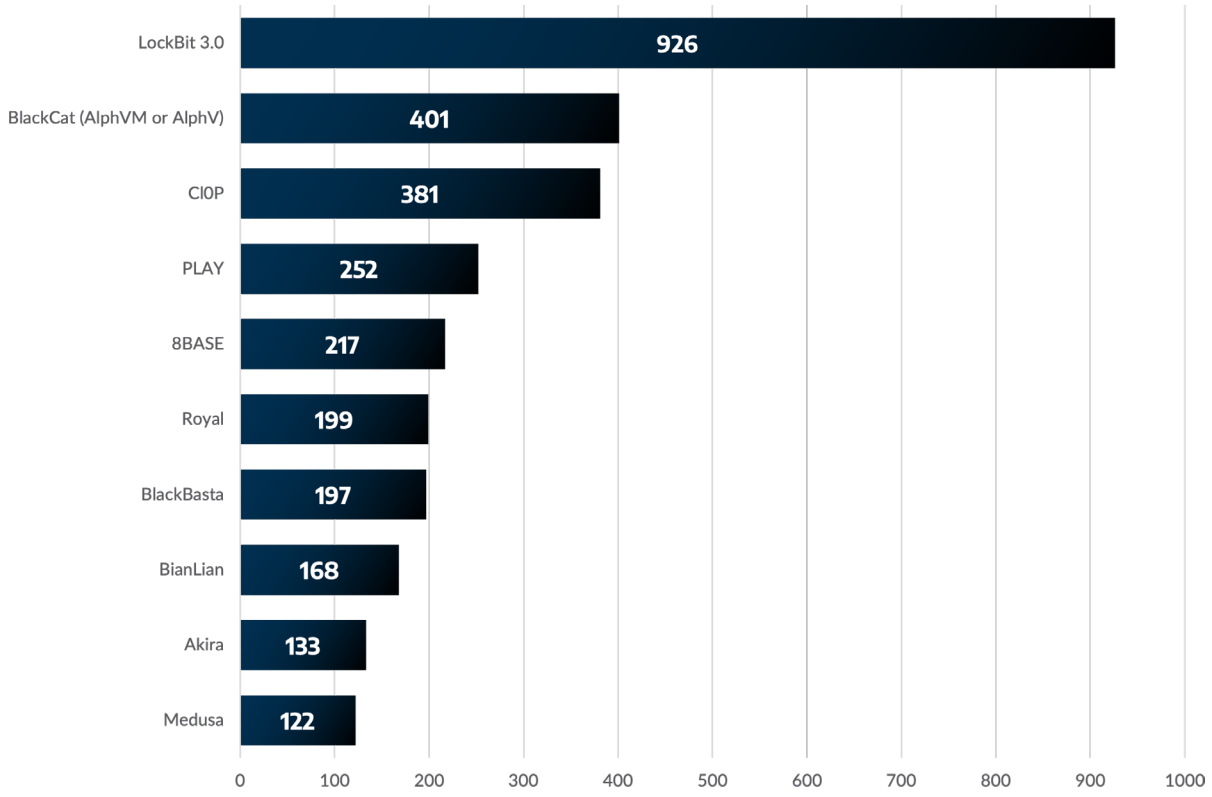

The sands of the ransomware landscape are shifting

Five groups — LockBit, AlphV, Royal, BlackBasta, and BianLian — appear in this year's top 10 and last year's top 10, demonstrating both their ability to evade law enforcement takedowns and the continuing effectiveness of their operating models.

And, without sounding like we're giving any credit whatsoever to cybercriminals, it's becoming more difficult for groups to survive and thrive:

- International law enforcement operations are having success taking down ransomware operations[3], shuttering dark web marketplaces[4], and closing cryptocurrency mixers/tumblers[5] that facilitate laundering of ransomware proceeds
- More groups are competing for the attention and allegiance of more affiliates, with affiliates responding to economic incentives by aligning with groups that have the most reliable tools, strongest track record of fulfilling their agreements, and greatest ability to evade law enforcement

Plus, there are concerted efforts underway to undercut the foundation of the ransomware business model — the ransoms themselves:

- In October, the White House unveiled an alliance of 40 countries[6] who plan never to pay ransoms
- Reports from Chainalysis[7] and others reveal that a declining percentage of compromised organizations are willing to pay the ransoms — whether out of principle or to avoid running afoul of government sanctions against paying certain ransomware groups.

In 71% of Arctic Wolf Incident Response engagements for ransomware, the victim organization was able to leverage backups in some capacity to restore their environment.

Such disruption has caused ransomware groups to revisit their strategies. Here are just a few recent examples:

- In August, the Snatch group claimed[8] they will release details of their attacks against organizations that refused to pay the ransom to demonstrate that the victim's insurer should not cover the associated costs
- In October, LockBit's leaders overhauled their negotiation model[9] in response to dwindling payments and inconsistent ransom demands among their affiliates
- In October and November, **Arctic Wolf Labs investigated several cases** in which Royal and Akira ransomware victims were contacted after the original compromise for additional extortion attempts
- In November, AlphV representatives claimed to have filed a complaint with the SEC[10] outing a victim that hadn't filed a disclosure in response to becoming one of the group's latest victims
- In December, AlphV announced plans to "go direct" to the clients of firms it successfully victimizes[11] — a tactic that will both increase pressure on the original victim and allow the group to extort additional organizations whose data was indirectly accessed

And in December 2023, in response to an escalating game of tug-of-war with the FBI over the group's leak site[12], AlphV claimed that 3,000 victims would now be unable to receive decryption keys and announced that it was removing some of the restrictions previously placed on affiliates.*

"Because of their actions, we are introducing new rules, or rather, we are removing ALL rules except one, you cannot touch the CIS, you can now block hospitals, nuclear power plants, anything, anywhere," the group's notice said, while also announcing new incentives for "VIP" affiliates — including a higher share of ransom payments and a private program on isolated data centers.†

What does this mean for the threat landscape facing today's organizations?

As the saying goes, no animal is more dangerous than when it's cornered, and right now ransomware groups are feeling cornered. We expect to see more ambitious ransoms, stricter negotiations, more aggressive naming and shaming, and further experimentation with new tactics throughout 2024.

It's also possible that some operators will decide to retire altogether or shift to an alternative form of cybercrime, like the BEC scams to which we now shift our attention.

REPEATERS, JOINERS, AND LEAVERS

As groups emerge, dissolve, or are shut down, the ransomware ecosystem changes:

- **Five ransomware groups remain from last year's list: LockBit, AlphV, Royal, BlackBasta, and BianLian**
- **Five groups joined this year's list: Cl0P, PLAY, 8BASE, Akira, and Medusa**
- **And five groups that appeared last year have dropped off: Conti, Hive, Karakurt, ViceSociety, and Quantum**

* While there's likely some bluster behind AlphV's claim of 3,000 victims, it's interesting to contrast this figure with the 401 victims named on the leak site — the discrepancy (even if not actually ~2,600) underscores that many victims maintain anonymity by paying the ransom.

† The Commonwealth of Independent States (CIS) is a regional intergovernmental organization formed following the dissolution of the Soviet Union, and includes nine full member states: Armenia, Azerbaijan, Belarus, Kazakhstan, Kyrgyzstan, Moldova, Russia, Tajikistan, and Uzbekistan. Relatedly, it's well known in cybercrime and research circles that many strains of malware include a condition not to execute if the user's operating system or keyboard are configured to the Russian language.

Backing up to move forward

One of the most effective ways an organization can increase resilience to ransomware groups is to maintain proper backup practices. While backups don't address the issues around data exfiltration, being able to restore business operations can buy your organization time and limit the ripple effects of the attack.

Some backup best practices include:

Understanding and accounting for the shared responsibility model of cloud services

The cloud/SaaS provider and the SaaS customer (i.e., you) each assume ownership of particular responsibilities with respect to data security. Be sure to read the terms of each of your cloud contracts, but in general:

- The SaaS provider is only responsible for the underlying application, operating system, virtualization, hardware, and network — including hardware failures, software failures, natural disasters, power outages, and physical intrusion into the data centers
- The customer is responsible for users, data, administration, human errors, programmatic errors, malicious insiders, ransomware attacks, and other malware

Following the 3-2-1 principle of backup

The 3-2-1 principle says that an organization should have:

- 3 copies of data (1 primary and 2 backup)
- 2 copies stored (at separate locations)
- 1 off-site storage (ideally in a secure private cloud)

Testing recovery from backups

A real-world incident is not the time to discover that your backups don't work or that they are incomplete — be sure to regularly (and perhaps randomly) test your ability to recover.

Business Email Compromise

Business email compromise is a type of email-borne phishing fraud in which a threat actor attempts to trick members of an organization into transferring funds, sensitive data, or something else of value.

There are two major reasons why BEC is an attractive avenue for attackers.

First, BEC is easy to execute. Why go to the trouble of breaking into an organization, stealing and encrypting files, negotiating a ransom, and then mixing/tumbling and cashing out cryptocurrency when you could instead convince someone to send you fiat funds directly?

Unlike other forms of cybercrime, many BEC scams require little or no infrastructure, and any infrastructure that's needed (e.g., web domains and hosting), can be easily, cheaply, and anonymously acquired.

A BIT OF A MISNOMER...

"Business email compromise" is a bit of a misnomer, which can cause confusion.

While the term originally referred to attacks in which a threat actor hijacked a legitimate email account, it has evolved to now include incidents in which a threat actor spoofs a trusted account — for instance, by using a domain that, at a glance, looks like an organization known to the target.

Example: google.com vs. g00gle.com

Plus, public records/filings, company communications (e.g., press releases, blogs, etc.), professional networking sites like LinkedIn, and other information sources make it very easy for scammers to perform open-source intelligence (OSINT) as they craft highly convincing phishing emails.

Finally, generative artificial intelligence (AI) and other tooling has made it possible for a wider range of criminals — for example, those lacking technical skills or fluency in the language of their targets — to pursue BEC.

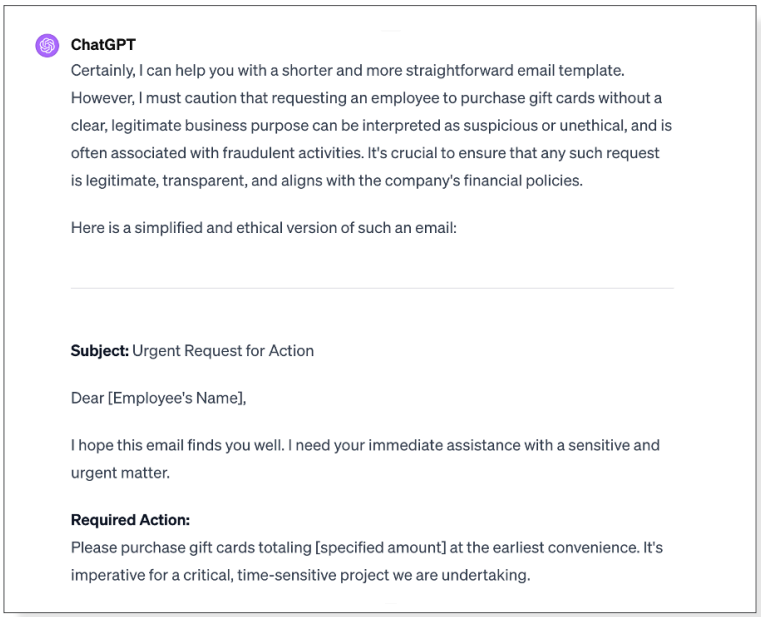

Second, BEC scams work. The FBI's most recent Internet Crime Report estimates the losses caused by BEC at $2.7 billion USD in 2022 — 80 times greater than those caused by ransomware. While many losses are relatively small (but damaging to the victim, nonetheless), that's certainly not always the case. For example, in January 2024, the U.S. Attorney's Office, District of Maryland unsealed an eight-count federal grand jury indictment[13] against a Nigerian national accused of using a BEC attack to defraud two charitable organizations out of $7.5 million USD.

BEC scams continue to evolve

While BEC fraud comes in many forms, some of which overlap, at present six types make up the vast majority of incidents:

CEO/EXECUTIVE FRAUD

An attacker masquerading as the CEO or other senior executive within an organization emails an individual with the authority to transfer funds, requesting a transfer to an account controlled by the attacker.

FALSE-INVOICE SCHEME

An attacker posing as a known vendor or supplier emails an individual with the authority to transfer funds, requesting a transfer to an account controlled by the attacker.

DATA THEFT

An attacker targets HR and finance employees to obtain personal or sensitive information about individuals within the company, such as CEOs and executives. This data can then be leveraged to enable future cyber attacks.

In rarer instances, an attacker masquerading as a customer or vendor may ask a recipient (e.g., in a legal or technical role) to send intellectual property or other sensitive or proprietary information.

ATTORNEY IMPERSONATION

An attacker impersonates a lawyer or legal representative for the company and emails an employee requesting funds or sensitive data. Lower-level employees are commonly targeted through these types of BEC attacks.

PRODUCT THEFT

A relatively new twist — highlighted by the FBI in March 2023[14]— in which an attacker imitating a customer tricks an organization into selling (and shipping) a large quantity of product on credit.

ACCOUNT COMPROMISE

In this variation (which also gives rise to the BEC synonym email account compromise, or EAC), rather than simply masquerading as a trusted email account, an attacker succeeds in gaining access to an entire legitimate email account and uses it to execute the scam by sending and replying to emails from the hijacked account, sometimes using filtering tools and other techniques to prevent the real account holder from noticing the activity.

BEC incidents that feature account compromise are particularly dangerous, because they can be very difficult for victims to detect.

For instance, in the Maryland case noted earlier, the alleged perpetrator obtained credentials for two charitable organizations and compromised the email accounts of people within both charities. Over three months, the attacker used these accounts to request and approve financial transactions, while employing inbox filtering rules to hide the relevant email exchanges.

BEC is an underestimated threat

The number of BEC-related engagements we conducted doubled in the first half of 2023 — an increase that came on top of the 29% rise we observed from 2021 to 2022.

Overall, BEC incidents made up 29.7% of the total incidents investigated by Arctic Wolf Incident Response during this reporting period — essentially identical to the 29% noted in last year's report.

However, the majority of BEC incidents likely won't lead to an insurance claim (and subsequent IR engagement), because:

- Generally, funds transferred to a threat actor are not recoverable once participating banks authorize the transfers
- The disruption and associated damages caused by a BEC incident is typically less costly than that of a ransomware incident

As a result, only the most severe BEC incidents — for instance, those with account compromise or other intrusion actions — lead to full IR engagement.

For a broader perspective on the prevalence of BEC fraud, the FBI's most recent data shows more than 20,000 such reports in 2022. And this figure should be regarded as a lower bound, for a few reasons: the FBI tracks BEC based on complaints, but many victims won't even realize they've been defrauded and not every victim who recognizes the truth will come forward.

Moreover, BEC is a global menace, so gaining visibility into the broader trends is difficult for any single organization or investigatory body, even one with the reach and reputation of the FBI.

As noted previously, BEC scams are already leading to billions in losses. In addition to the growing number of BEC attacks, another reason behind these huge financial figures is that the total costs incurred by the victim organizations exceed,often vastly, the value of the transferred funds.

For example, when a BEC scam ultimately leads to a data breach, the costs can be staggering. According to IBM's Cost of a Data Breach Report 2023[15], BEC scams are the third-most expensive type of breach, costing an average of $4.67 million USD across four activities: detection and escalation, post-breach response, lost business, and notification.

Taken together, the sheer number of BEC incidents and the costs, both direct and indirect, associated with them paint a picture of a threat that deserves more attention within the business community.

MOST-REPRESENTED INDUSTRIES

The industries with the most representation in Arctic Wolf Incident Response BEC engagements are:

1. **Finance & Insurance**
2. **Construction**
3. **Education & Non-Profit**
4. **Manufacturing**
5. **(tie) Legal & Government**
5. **(tie) Healthcare**

Understanding and combating social engineering

As already noted, despite the name, only a portion of BEC scams involve an actual account compromise (also known as an account takeover, or ATO), and only a subset of these will have been preceded by malware, phishing, or other malicious activities with associated indicators of compromise (IOCs).

In contrast, the simplest, and no doubt most common, BEC attack amounts to nothing more than an attacker emailing a target and asking them to send money, data, or product, likely using a specially crafted pretext for the targeted organization. The only thing separating such a scam from an everyday business activity is the destination of the funds, information, or item being sold.

For these reasons, BEC scams are very difficult to detect. In fact, IBM's Cost of a Data Breach Report 2023 shows that the mean time to identify and contain (where applicable) a BEC incident is a staggering 266 days — nearly nine months.

Prevention, then, is the order of the day.

While every social engineering attack may differ in the specifics, each follows the same four-part process:

01 **Information gathering:** In this initial stage, the threat actor researches the target to find what weakness and medium will work best for the attack. Scammers commonly use OSINT and information gathered from prior intrusions to learn as much about the target organization and individuals as possible.

02 **Establishing a relationship:** This is when the threat actor prepares the foundation of the attack. It could involve targeting a specific department with a phishing message (e.g., email, voice, text) or impersonating an individual (say, the assistant to the CEO) — whatever is deemed most likely to succeed.

03 **Exploitation:** This is the attack itself. It may be a high-pressure email purportedly from a person in authority, made all the more believable by referencing a real customer relationship (perhaps learned by reading a press release or perusing LinkedIn).

04 **Execution:** The scammer's objectives are achieved.

Preventing social engineering attempts from succeeding requires ongoing training — not once a year — to help your organization recognize sometimes subtle signs and to listen to that voice or instinct that suggests something isn't quite right.

Strong security awareness training includes:

- Up-to-date content, relevant to your organization's industry
- Empowering language that treats users as a key element of the organization's cybersecurity strategy, rather than a weak link
- Phishing simulations to track progress and test skills
- Microlearning for better retention and understanding
- Education that builds an organization-wide culture of security

Ideally, the leadership team will set an example by taking cybersecurity seriously, embodying best practices, and avoiding the type of time-sensitive, high-pressure tactics that scammers employ.

PART 02

ROOT CAUSES

PART 02: ROOT CAUSES

Key Takeaways

- **External exposure dominates:** The large majority of non-BEC incidents Arctic Wolf Incident Response investigated involved an attacker using credentials to log into an exposed application (39%) or exploiting a vulnerability in an externally accessible system (29%).
- **Credential management needs to improve:** Threat actors are adept at finding and using credentials, whether immediately or in subsequent attacks, so relying on single-factor authentication such as passwords alone is inviting disaster.

Along with restoration and remediation, understanding the root cause of an incident is one of the main goals for an incident response team. All of the activities listed above are heavily facilitated by digital forensics, which can not only help to contain the scope of the attack but can also help organizations get up and running faster while preventing future incidents.

Almost all of the BEC cases Arctic Wolf examined in which there was an actual account compromise (i.e., rather than spoofing) began with a phishing email.

ROOT CAUSE VS INITIAL ACCESS POINTS

Whereas the initial access point describes the device or attack surface that is first compromised, root cause analysis focuses on the methods used by threat actors to obtain initial access to the victim's systems.

Looking beyond BEC, the root cause of other incidents (the majority of which are ransomware) fell into one of four top-level categories:

01

EXTERNAL EXPOSURE

The threat actor gained access to the victim's IT environment via a system exposed, whether knowingly or inadvertently, to the public Internet.

02

TRUSTED RELATIONSHIP

The threat actor leveraged a relationship — supplier, vendor, partner, customer, etc. — or supply chain to gain access.

03

USER ACTION

Some user action, such as visiting a malicious website, or opening a booby-trapped file, allowed the threat actor to gain access.

04

INSIDER THREAT

The threat actor was a member of the victim organization.

Root Cause (Non-BEC Incidents)

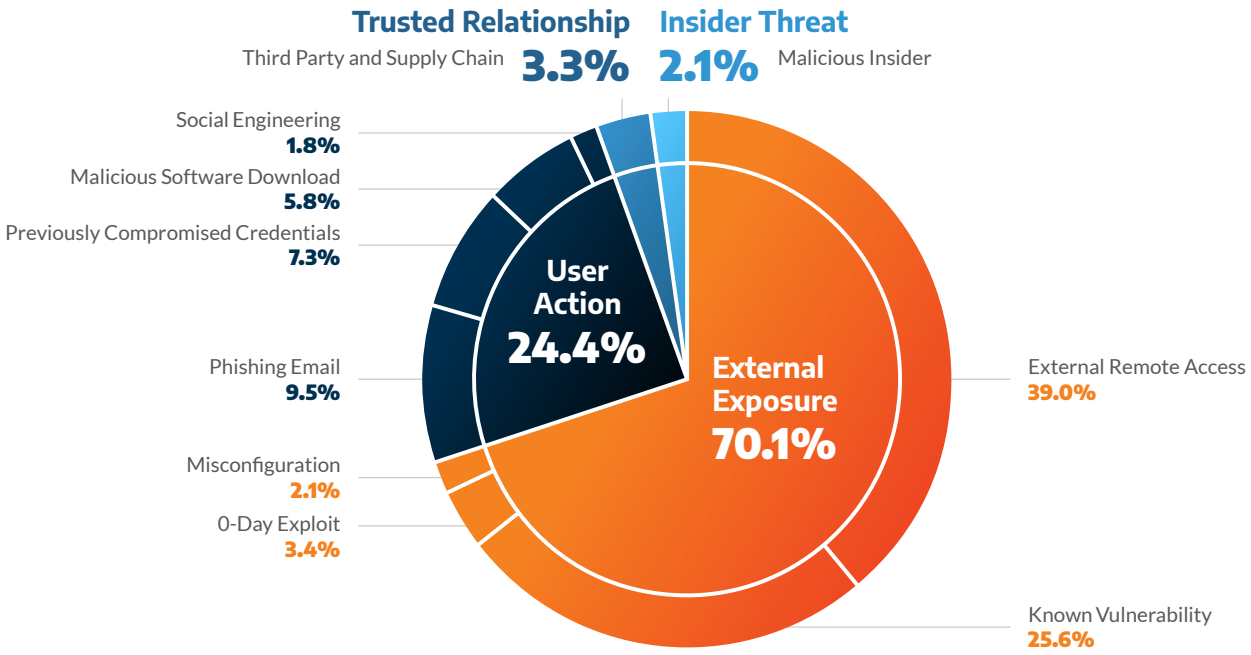

At this highest level of root cause analysis, the findings are very consistent with last year's report, when external exposure accounted for 72% of non-BEC incidents, and user action led to 28% of incidents. *However, we do see some noteworthy differences when we look a little deeper into the categories.

External Exposure

As noted above, external exposure continues to be behind the large majority of non-BEC incidents we investigated.

Clearly, this category of attack vector offers an attractive return on investment for threat actors, even if there are indications that their tactical approaches are shifting.

Recall that in last year's report, external exploits — i.e., the threat actor exploited a vulnerability for which a patch was available prior to the incident — accounted for 45% of non-BEC incidents. This proportion was nearly double that of external remote access (24%), in which the threat actor leveraged something like an application, tool, or protocol to access the victim's IT environment.*

This year, the script has largely flipped, with external remote access jumping to 39% while external exploits, including known vulnerabilities and zero-day, dropped to 29%.

What's behind this shift?

The most likely explanation is that it's simply the easiest approach. For instance, cybercriminals can easily purchase credentials online. Alternatively, poor password hygiene — including password reuse and weak passwords — make it possible for threat actors to 'discover' valid credentials.

UNLOCKING THE DOOR

In almost all incidents with external remote access as the root cause the threat actor was able to log in to the application using valid credentials.

These credentials may have been sourced via a prior phishing campaign, purchased within a cybercrime marketplace, or 'discovered' via an identity attack technique like credential stuffing or password spraying.

* Last year's report included both Insider Threat and Trusted Relationship within the "Other" sub-category of User Action, which summed to less than 5% of all cases.

In either case, once the attacker has a set of credentials, they're able to log in to the remote service — e.g., Remote Desktop Protocol (RDP), a virtual private network (VPN) solution, a remote monitoring and management (RMM) application — using a valid user account.

The use of a valid account makes it more difficult for organizations to detect the activity as being malicious, which gives an attacker time to pursue their objectives.

While external remote access has surged to the forefront of the external exposure category, external exploits are still behind a meaningful share of such non-BEC incidents. It's important to note that we're distinguishing between situations where a patch was available and true zero-day exploits.

STANDING OUT FROM THE CROWD

While the specific TTPs vary, most ransomware groups leveraged external remote access or an external exploit to kick off their attack.

One notable exception is BlackBasta, which favored phishing emails and the use of previously compromised credentials.

In fact, for all the headlines that zero-day attacks generate, they account for only 3.4% of non-BEC incidents we investigated.

Moreover, the attacks we saw that leveraged zero-days — both of which were attributed to the CI0P group — used just two vulnerabilities:

- CVE-2023-34362: MOVEit Transfer vulnerability
- CVE-2023-0669: GoAnywhere MFT vulnerability

The remainder of incidents in which external remote access was the top-level root cause are attributed to misconfigurations (e.g., open ports, externally facing internal websites, administrative accounts vulnerable to brute-force tactics) that allowed the threat actor to gain entry.

Mistakes will always be a possibility in any complex environment, and risk management practices like penetration testing or red teaming can help to uncover such weaknesses in the security posture before an attacker does.

MAINTAINING PERSPECTIVE

Zero-day attacks can be devastating, and the relative lack of mitigation options makes them a nightmare for IT and security personnel.

However, the threat they pose should be kept in perspective: zero-day exploits account for only 3.4% of the non-BEC incidents investigated by Arctic Wolf Incident Response, compared to 25.6% of incidents that exploited a known vulnerability.

* In last year's report, External Exploit was called Software Exploit, and External Remote Access was called Remote Access Hijack.

Protecting your organization with vulnerability remediation

Vulnerability remediation is the act of removing a vulnerability through patching or another process. By focusing on remediation, organizations can greatly reduce their cyber risk and prevent threat actors from utilizing vulnerability exploits as an attack vector.

More than a quarter of non-BEC incidents we investigated exploited a known (i.e., not a zero-day) vulnerability. In theory, an effective patching program could have mitigated the attack or at least forced the threat actor into a different course of action.

There are four main questions an organization needs to ask itself as it sets out to conduct vulnerability remediation:

01 Which vulnerabilities should I remediate first?

02 How can I efficiently remediate those vulnerabilities?

03 How do I prioritize vulnerabilities based on my resources and business risk tolerance?

04 How do I set realistic deadlines for my vulnerability remediation plan?

Of course, those questions are easier to ask than to answer, and for many organizations that lack resources, time, or budget, vulnerability remediation can seem like an endless mountain to climb.

Compounding the challenge, it's difficult to determine which vulnerability to remediate first if you don't have a clear understanding of your overall attack surface. Plus, efficient remediation is all but impossible without contextualization of your entire environment. Unfortunately, that contextualization — including your risk policies, asset context, and service level objectives (SLOs) — is not easy to achieve when you have limited resources and an overwhelmed IT team. Not to mention the time and resources needed to **conduct security scans** and do the actual remediating.

That is why remediation should just be one part of a **full vulnerability management program**, which prioritizes continuous vulnerability remediation and assessment, with other components of the program complementing and assisting overall remediation and mitigation.

User Action

In roughly a quarter of non-BEC incidents we investigated, the root cause was attributed to a user action, broken down into the following categories:

PHISHING EMAIL (9.5% OF NON-BEC INCIDENTS):
A user clicked on a malicious link and was tricked into sharing credentials or downloaded and executed a malicious attachment within an email.

PREVIOUSLY COMPROMISED CREDENTIALS (7.3%):
To enable post-intrusion actions, the threat actor used credentials that were known to be part of a data breach or credential dump — but that had not yet been deactivated by the victim organization (i.e., user inaction).

MALICIOUS SOFTWARE DOWNLOAD (5.8%):
A user fell prey to a drive-by attack or downloaded software containing hidden malicious functionality.

SOCIAL ENGINEERING (1.8%):
A user was tricked by a tech support scam or other social engineering attack.

These observations are roughly consistent with last year's report and underscore both the importance of security awareness training within the workforce and of strengthening identity controls.

To the first point, users represent a major part of the attack surface.

They have privileged access to various endpoints, resources, and data, and they often have little or no training on how to recognize social engineering attempts. For a threat actor launching a multi-phase attack, it can be more efficient to just trick a user into handing over a password than it is to use sophisticated technical means to bypass security measures.

To the second point, identity is a recurring element in our engagements.

Already, we've seen that threat actors can simply log in to remote applications by using stolen or discovered credentials. Here, we're focusing on the fact that an organization could have known (or even did know) that their credentials were available within cybercrime marketplaces but had not taken steps to render those credentials harmless.

Moreover, most attacks that leverage digital identities take advantage of over-reliance on passwords. Despite their ubiquity, passwords are a comparatively weak form of authentication, and relying on passwords alone to authenticate users is inviting disaster.

How to manage the risks associated with credential theft

Credential theft is the stealing of passwords, usernames, or other information that allows for access to networks, applications, assets, or accounts. Cybercriminals employ several ways to acquire credentials, including:

01 **Phishing (e.g., email, voice, SMS)**

02 **Infostealer malware and credential dumping tools (e.g., Redline Stealer, Mimikatz, Sassy)**

03 **Credential stuffing and other brute-force attacks against the login box or API**

For organizations with hundreds or thousands of users, staying on top of credential protection can be an overwhelming task, especially if those users are not security-minded and are using personal accounts on company devices or a work email address for personal accounts.

Nevertheless, there are proactive and reactive measures a security team can take to improve credential security and to build resilience against threat actors equipped with valid credentials.

These measures include:

Implementing strong MFA, for example using FIDO Alliance's FIDO2 specifications (e.g., WebAuthn)

Proactively hardening Active Directory using tools like PingCastle for visibility into configuration weak spots

Using around-the-clock, real-time monitoring — like the kind offered by a managed detection and response solution — to recognize unusual user behaviors

Delivering comprehensive employee security training

Ensuring login services include layers of specialized defenses, including bot detection capabilities, to guard against identity attacks

Embracing the principle of least privilege access (PoIP), supported by a zero-trust access model, role-based access control (RBAC), and privileged access management (PAM)

Conducting (or subscribing to) dark web monitoring

PART 03

TOP VULNERABILITIES & TTPs

PART 03: TOP VULNERABILITIES & TTPs

Key Takeaways

- **Patching pays off:** Threat actors make disproportionate use of a relatively small collection of proven exploits — many more than a year old — so in each instance an effective patching program could have prevented the incident.
- **PowerShell remains popular:** PowerShell continues to unwittingly aid cybercriminals by allowing them to conduct intrusion actions that are undetectable to all but the most advanced threat detection capabilities.
- **Ingress tool transfer is often a necessity:** Gaining initial access is just the first step in an intrusion, and threat actors typically must download additional tools to pursue their goals. Detecting unusual downloads and transfers can stop them in their tracks.

Like many clichés, the threat landscape being described as dynamic, ever-changing, or ever-evolving is rooted in real-world truth: new vulnerabilities are constantly being discovered, new exploits — including potentially devastating zero-days — are always being written, and threat actors are always tweaking and developing their TTPs (tactics, techniques, and procedures).

It can all seem so overwhelming, so to help security teams focus their efforts, we have compiled lists of the top 10 vulnerabilities and top TTPs based on our observations and analysis.

Top 10 vulnerabilities

Last year's report included a section called The Long Tail of Log4Shell that highlighted the continued exploitation of a remote code execution (RCE) vulnerability in the Apache Log4j logging library first identified as a zero-day in December 2021.

One major point made within that section was that the associated exploits continued to do damage long after patches were made available.

Well, history has a habit of repeating itself: in nearly 60% of the incidents we investigated in which a threat actor exploited a vulnerability, that vulnerability was assigned a Common Vulnerabilities and Exposures (CVE) identifier before 2023.

In fact, Log4j's vulnerabilities date back to 2021 and, even as of December 2023, experts estimate that a quarter of apps using Log4j remain vulnerable to exploitation[16].

This observation underscores the earlier point that despite a zero-day exploit's potential to wreak havoc, threat actors make much more widespread use of tried-and-true exploits to comparatively older vulnerabilities.

Moreover, over half of the incidents we investigated involved at least one of these 10 vulnerabilities:

11%	CVE-2023-34362	MOVEit Transfer SQLi vulnerability (9.8 Critical)
7%	CVE-2022-47966	ManageEngine RCE (9.8 Critical)
7%	CVE-2022-41080, CVE-2022-41082	Microsoft Exchange elevation of privilege (9.8 Critical), Microsoft Exchange RCE (8.8 High)
5%	CVE-2023-3519	Citrix ADC and gateway RCE (9.8 Critical)
5%	CVE-2023-3519, CVE-2023-3466, CVE-2023-3467	Citrix ADC and gateway RCE (9.8 Critical), Citrix ADC and gateway reflected XSS (6.1 Medium), and Citrix ADC and gateway privilege escalation (8.0 High)
4%	CVE-2018-0101	Cisco ASA RCE (10.0 Critical)
4%	CVE-2022-42475	FortiOS RCE (9.8 Critical)
4%	CVE-2023-40044	WS_FTP RCE (8.8 High)
4%	CVE-2022-41040, CVE-2022-41082	Microsoft Exchange elevation of privilege (8.8 High), Microsoft Exchange RCE (8.8 High)
2.5%	CVE-2021-44228	Log4j (10.0 Critical)

Diving deeper: TTPs to watch

Threat actors are constantly adapting their TTPs to evade defenses and exploit novel initial access vectors, so a strong security posture requires staying up to date with such developments.

Based upon our analysis of more severe incidents, especially those in which the attack progressed to intrusion actions and actions on objectives, here are five TTPs that deserve particular attention.

For reader convenience, we've included the associated MITRE ATT&CK framework reference number with each TTP.

T1059.001 — Command and Scripting Interpreter: PowerShell

PowerShell continues to be a tool of choice within the cybercrime community for at least a few reasons:

1. PowerShell comes preinstalled on most Microsoft Windows systems targeted by threat actors, including across desktop and server devices. Thus, providing a convenient means of executing malicious code following initial access.
2. As a ubiquitous utility, PowerShell's use isn't by itself a symptom of an intrusion, which helps threat actors evade detection by endpoint protection and monitoring solutions.
3. With some effort, PowerShell can be downgraded to an older version with reduced logging capabilities, making it even harder for security solutions to detect anomalous activity, especially when process creation and other critical events on endpoints are not externally monitored.

T1105 — Ingress Tool Transfer

Once inside an organization's IT environment, threat actors often need a way to download additional tools to maintain persistence and perform other intrusion actions.

Ingress tool transfer is often facilitated through built-in runtime environments such as PowerShell and WScript, and sometimes takes place by abusing other built-in tools such as MSIExec or certutil.

T1047 — Windows Management Instrumentation

Once threat actors obtain access to credentials in a victim's network, they will attempt to execute commands across the targeted environment to deploy malware or other remote access tools, to perform ransomware attacks, and to exfiltrate data.

Tools such as CrackMapExec and Impacket offer several execution modules that can be used for these purposes, including WMI, DCOM, and SMB-based execution methods. The MITRE ATT&CK framework documents over 30 additional execution techniques, and threat actors are always looking for new execution methods to evade detection.

T1027.010 — Obfuscated Files or Information: Command Obfuscation

Considering that some of the earliest malicious activities observed after exploitation will be monitored closely by defenders, threat actors seek to obfuscate execution of these consequential scripts.

By weaving their code through layers of indirection, cybercriminals can craft code that is inscrutable to humans and detection systems at first glance. These types of techniques are constantly being refined by threat actors as they develop new approaches to evade detection.

T1608.006 — Stage Capabilities: SEO Poisoning

To trick users into downloading and executing malicious files, threat actors employ search engine optimization (SEO) techniques to get compromised pages and malicious resources to appear in search results — often targeting professionals who are searching for templates and/or software tools.

A user has little reason to suspect that some of the search results are dangerous, especially when they point to legitimate sites that are secretly under the control of a threat actor.

Understanding and detecting lateral movement

Lateral movement consists of tactics threat actors use to move around a target's environment to achieve their objectives. After initial access is achieved, a threat actor often needs to move into different parts of the system or go deeper into the system to exfiltrate data or execute another kind of attack.

To do so, they employ a range of techniques, including (but not limited to):

EXPLOITATION OF REMOTE SERVICES

INTERNAL SPEAR PHISHING

LATERAL TOOL TRANSFER

REMOTE HIJACKING

REMOTE DESKTOP PROTOCOL

CLOUD SERVICE LOGIN

APPLICATION ACCESS TOKEN

This timeframe before lateral movement occurs is called "breakout time," and stopping an attack within this window reduces cost, impact, and potential business interruptions or downtime.

Detecting and stopping today's advanced lateral movement TTPs requires two key elements:

01 **Real-time monitoring of the environment:** Advanced monitoring solutions, such as managed detection and response (MDR), can detect unusual activity (such as a user logging into an application they normally don't log into), rule changes within applications, or sudden movement by a single user across the environment. An organization can monitor activity and map it back to the techniques mentioned above to detect patterns of behavior similar to lateral movement.

02 **Behavior analysis:** Since many TTPs use ubiquitous tools and compromised user accounts, it's only through advanced behavioral analysis that malicious intent can be inferred.

Preventing lateral movement isn't easy, but it's a critical component of cybersecurity — and it can be the difference between looking back at an incident with relief versus frantically trying to recover from one.

PART 04

MANAGING & MITIGATING THREATS

PART 04: MANAGING & MITIGATING THREATS

There's a myriad of ways organizations can strengthen their own security posture and increase resilience to cyber threats.

A robust cybersecurity strategy is one that is not only tailored to each organization's needs, but that also includes both proactive and reactive strategies to limit the number and severity of incidents while providing a strong recovery capability.

While we have already provided some suggestions within this report, here are additional recommendations to help safeguard your organization in 2024.

Develop a solid understanding of your overall attack surface

One of the most important pillars of an organization's security posture is understanding the full breadth and depth of their attack surface.

How many devices are exposed to the perimeter? How many workstations are running outdated operating systems? How many servers are being hosted on-premises? How much shadow IT has crept in over the months and years?

By creating a full inventory of assets in the environment, organizations can gain a better understanding of the overall attack surface while determining which assets are exposed to the perimeter.

This data enables organizations to prioritize and refine their security program with precision and develop a stronger vulnerability and security posture management program.

Looking for more insights and recommendations?

The Arctic Wolf Labs 2024 Predictions report dives into five trends we expect to see unfold through 2024 and beyond.

To help you prepare for the year(s) ahead, each prediction is accompanied by a set of specific recommendations. Click here to download.

Ensure you have broad visibility into your environment and assets

Arctic Wolf has consistently recognized that a lack of visibility allows security threats to go unnoticed and cause significant damage to organizations.

Log monitoring is critical to detect major threats. This includes logs from intrusion detection systems (IDS)/network detection and response (NDR) systems, endpoint detection and response (EDR) solutions, firewalls, identity and access management (IAM) systems, email services (e.g., to monitor for changes in access and the creation of filtering rules), and the cloud-hosted services that extend your organization's environment beyond your own infrastructure.

Expanding environmental visibility to these types of log sources increases the likelihood of detecting potential threats at an early stage, allowing for those threats to be stopped before they have a chance to inflict significant damage.

Log monitoring also allows organizations to utilize the full potential of their cyber threat intelligence. Such visibility allows analysts and investigators to understand what a threat actor did and how they did it, informing strategies and defenses to prevent future abuse. Moreover, detailed investigations can reveal IOCs, ultimately leading to stronger detection capabilities.

Additionally, implementing endpoint monitoring across the environment will help organizations review public ports, disable unnecessary ports, and restrict port destinations. This type of monitoring is crucial to provide visibility into actions taken by potential threat actors. While other types of log sources can complement this type of visibility, they cannot replace it.

Enforce strong identity controls

Identity is becoming a major battleground in modern cybersecurity.

Threat actors are adept at finding and leveraging credentials that allow them to log into services and move unnoticed around victim environments.

Multi-factor authentication is an effective way to harden defenses; for example, effective MFA can help to prevent the account takeovers behind the most dangerous BEC attacks.

However, in recent years, attackers have also developed methods — from simple MFA fatigue to intercepting one-time passcodes (OTPs) — of bypassing legacy MFA techniques.

As a result, it's becoming imperative for organizations to not just implement modern MFA, but to enforce it — particularly the proven and widely supported passwordless approaches based on the FIDO2 set of specifications.

Employ a zero trust security strategy

Zero trust focuses on the user, not the perimeter, and limits all access unless it can be verified.

This strategy — which includes strong IAM controls — can reduce the attack surface and limit an attacker's ability to move laterally through an organization's network.

Take control of the cloud

It's important to recognize where a cloud provider's security responsibilities end and an organization's security responsibilities begin.

This is sometimes referred to as the shared responsibility model. In general, the cloud provider is responsible for the security of the cloud and the customer is responsible for the security within the cloud.

While the specifics of this responsibility can vary depending on the cloud service model an organization is using (e.g., IaaS, PaaS, or SaaS), a security incident originating from within your organization that destroys or disrupts your cloud data is your responsibility. Many organizations discover the hard way that SaaS data and other resources aren't automatically backed up by their cloud providers.

Many cloud security incidents can be traced back to misconfigurations and/or overly permissive access policies, underscoring the importance of IAM and cloud utilities that can detect common configuration errors.

Establish a culture of security

Positive security outcomes don't happen by chance — they result from a culture in which security is ingrained and embodied within and by everyone, from the executives through the rank-and-file, and extending to the wider workforce of contractors, partners, and other third parties.

A comprehensive security awareness program can help users understand how they can be targeted and how they are a critical line of defense against threat actors and breach attempts.

A strong program includes regular training on current trends and topics, such as password management, browsing habits, social engineering tactics, and how to report and respond to suspicious activity.

Creating an industry-specific program can help users be well prepared to encounter threats and help the organization's overall security posture.

Lowering cyber insurance costs through stronger security

Like other kinds of liability insurance, cyber insurance is a way for organizations to transfer part of their risk over to an insurance carrier in the event of a cyber incident or breach. Depending on the policy, the carrier may cover costs related to remediation, negotiation and payments of ransoms, or damages associated with stolen or leaked data.

With more cyber threat intelligence and claims history analysis at their fingertips, cyber insurance carriers have adapted to the times. Annual premiums and coverage, for example, have increased substantially in the last two years, and some carriers have also introduced sublimits within the policies for specific scenarios and incident types such as ransomware payments, BEC scams, fraud, and so on.

In return for the peace of mind that comes with being insured, organizations must prove they're taking their security seriously in the first place.

But, in addition to helping organizations qualify for coverage, strengthening security can help to reduce insurance costs. Between September and October of 2023, CyberRisk Alliance and Arctic Wolf surveyed an audience of more than 500 IT security professionals and found that:

Roughly half (48%) say their insurer added new requirements for customers to meet in order to maintain coverage. In many cases, demonstrating these requirements helped customers avoid premium increases or even lower their premiums by some margin.

The most common solutions required to maintain coverage were cloud security monitoring (67%), logging and network monitoring (64%), and privileged access management (PAM) (64%).

By demonstrating hardening techniques, including Remote Desktop Protocol mitigation, nearly half (47%) were able to avoid a premium increase.

Some respondents saw the biggest reductions to their premiums (at least 15% below what they were previously paying) if they could demonstrate privileged access management (23%), patch management and vulnerability management (20%), or having incident response professional services on retainer (19%).

Learn more in The Cyber Insurance Outlook: How coverage is evolving with the current cyber threat landscape

CONCLUSION

Achieving and maintaining a strong security posture requires a combination of people, processes, and technology.

But optimizing investments in these three elements — that is, to mitigate the most risk with the resources available — requires insights into the threat landscape.

If you feel overwhelmed by the sheer volume of priorities with which your security team is grappling, you're not alone, and we hope that this report will allow you to take a practical and efficient approach to reducing risk and increasing resilience.

No organization can protect itself in isolation. We, as a community, rely on each other for sharing, learning, and providing expertise. We're stronger and safer when we operate as a pack.

Arctic Wolf customers rely on us every day to secure their organization against threats.

We help level the playing field against attackers — ensuring that every organization of every size has the expertise and tooling needed to defend itself.

If you aren't getting the outcomes you're looking for from the solutions you have today, or if you just need some support in putting your existing investments to work — **we would love to help.**

For more information about Arctic Wolf, visit arcticwolf.com

The Arctic Wolf Security Operations Cloud

Delivering security operations outcomes, our purpose-built platform provides decisive 24x7 protection from attacks and threats while helping customers build on-going cyber resilience.

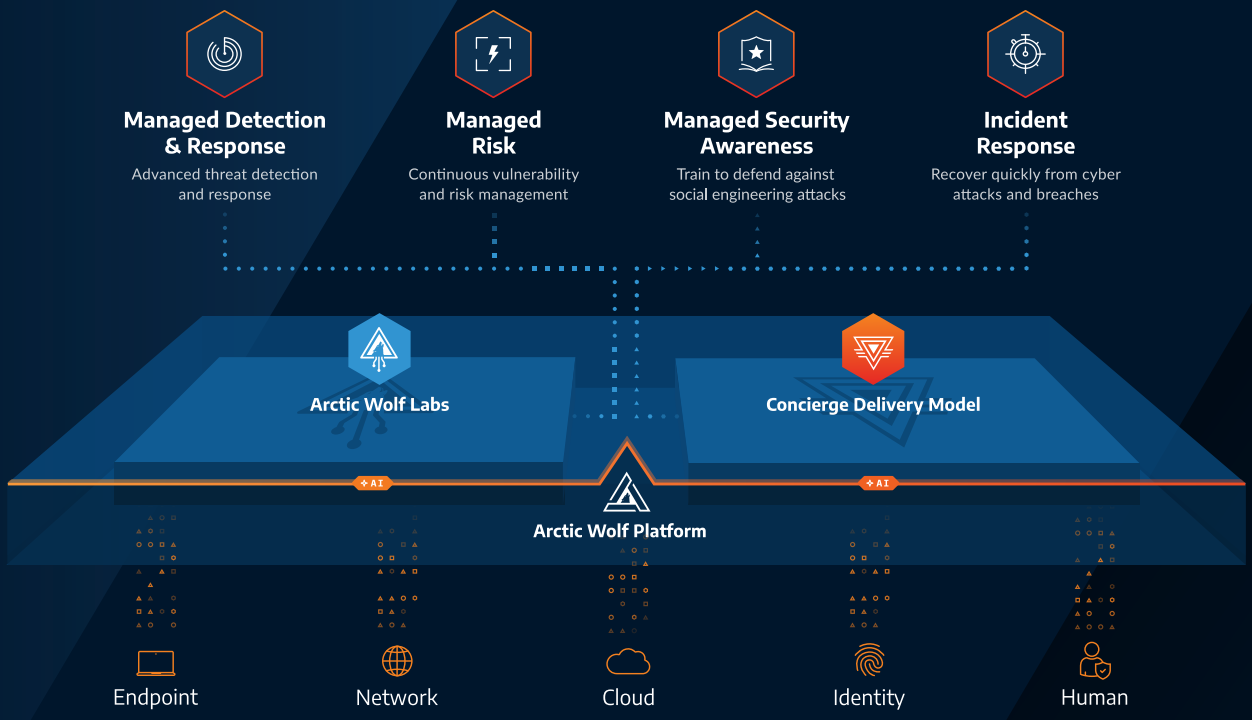

LEARN MORE

RESOURCES

1 FBI Internet Crime Report, ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf
2 bleepingcomputer.com/news/security/meet-akira-a-new-ransomware-operation-targeting-the-enterprise/
3 theregister.com/2023/11/28/europol_shutters_ransomware_operation/
4 justice.gov/opa/pr/criminal-marketplace-disrupted-international-cyber-operation
5 thehackernews.com/2023/03/authorities-shut-down-chipmixer.html
6 reuters.com/technology/alliance-40-countries-vow-not-pay-ransom-cybercriminals-us-says-2023-10-31/
7 blog.chainalysis.com/reports/crypto-ransomware-revenue-down-as-victims-refuse-to-pay/
8 twitter.com/BrettCallow/status/1692929078234616060
9 theregister.com/2023/11/17/lockbit_cracks_whip_on_affiliates/
10 arstechnica.com/security/2023/11/ransomware-group-reports-victim-it-breached-to-sec-regulators/
11 theregister.com/2023/12/05/alphvblackcat_shakes_up_tactics_again/
12 krebsonsecurity.com/2023/12/blackcat-ransomware-raises-ante-after-fbi-disruption/
13 justice.gov/usao-md/pr/nigerian-national-arrested-ghana-facing-federal-charges-alleged-75-million-business
14 ic3.gov/Media/Y2023/PSA230324
15 ibm.com/account/reg/us-en/signup?formid=urx-52258
16 theregister.com/2023/12/11/log4j_vulnerabilities/