Mandiant

Google Cloud Security

M-Trends

2025 Report

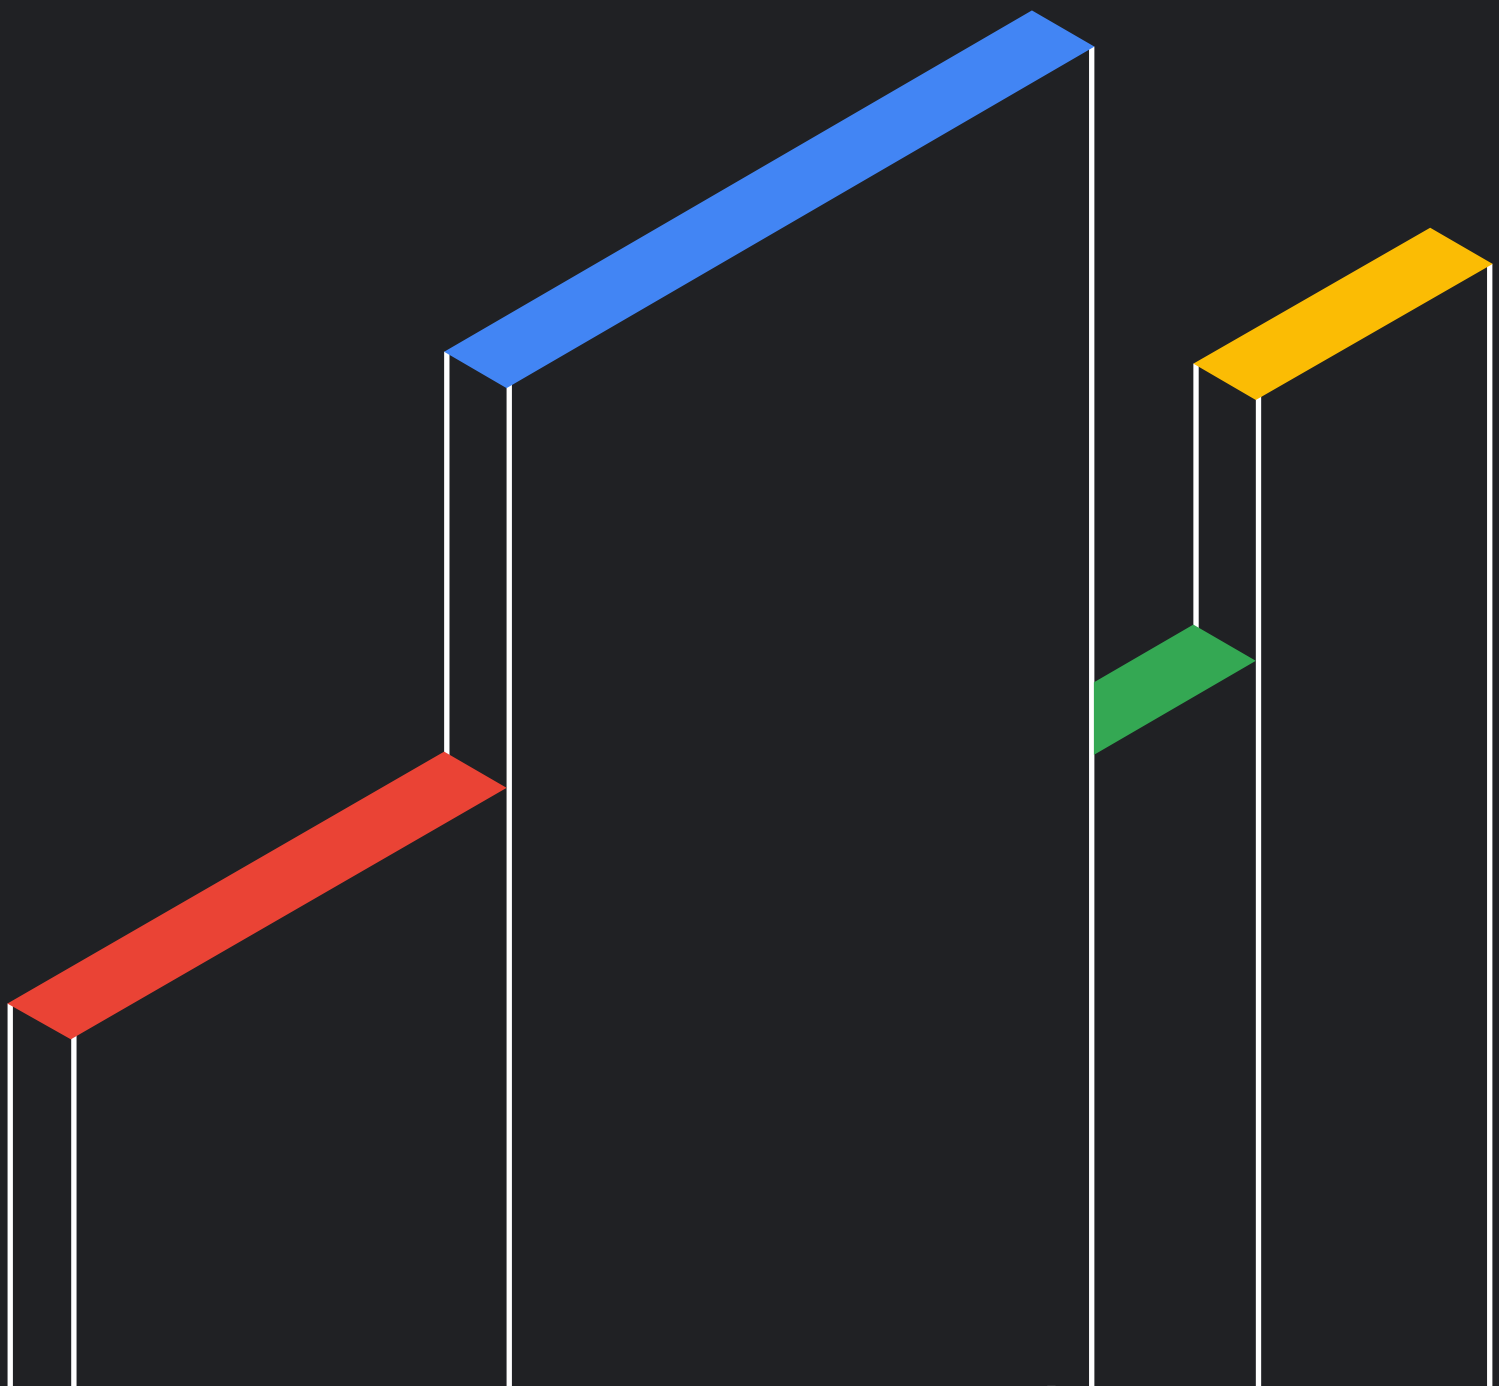

Table of Contents

Introduction

A key takeaway from M-Trends 2025 is that attackers are seizing every opportunity to further their objectives. One way they are doing this is through the use of infostealer malware, which is increasingly being used to enable intrusions using stolen credentials. Another growing trend is the targeting of unsecured data repositories, which is brought on by the lack of basic security hygiene. Additionally, attackers are exploiting the gaps and risks introduced as organizations continue their migrations to the cloud.

The most common way attackers breached organizations in 2024 was through exploits, which we observed as the initial infection vector in 33% of our investigations. The financial sector continues to be the most targeted industry, making up a little more than 17% of our investigations. Global median dwell time has risen to 11 days from 10 days in 2023. This marks the first increase since the publication of the inaugural M-Trends in 2010 but is still below the 16 days reported in 2022. In M-Trends 2025, we take a look at how adversary notifications—notably in ransomware incidents—influence the global median dwell time metric.

By providing data and other security metrics in M-Trends, along with deeper dives on attacker trends, we illustrate how threat actors are conducting their operations, how they are achieving their goals, and what organizations need to be doing to prevent, detect, and respond to threats. Infostealer malware, unsecured data repositories, and cloud migrations are just a few challenges organizations will face. We additionally cover:

- Insider risk brought on by Democratic People's Republic of Korea (DPRK) IT workers
- Growth of blockchain technology leading to cryptocurrency and Web3 threats
- Iran-nexus threat actor operations amid Middle East tensions

Mandiant consultants are regularly on the frontlines of cyber incidents, where they conduct in-depth investigations and analysis of the most recent attacks. This firsthand experience results in a deep understanding of threats and the effective strategies required to defend against them.

Mandiant uses this knowledge to proactively assess client security postures, comparing them against the latest attacker tactics, techniques, and procedures. Furthermore, we provide critical support for remediation efforts, security transformation initiatives, and comprehensive security education.

Through the release of our annual M-Trends report, we share our learnings with the greater security community, building on our dedication to providing critical knowledge to those tasked with defending organizations. The information in this report has been sanitized to protect the identities of victims and their data.

By the Numbers

Since 2010, Mandiant has provided statistics and analysis of threats observed in the previous year's incident response investigations. In M-Trends 2025, Mandiant examines data collected from more than 450k+ hours of incident response engagements globally, highlighting trends and significant insights. This information can be useful to inform risk assessments and to support planning for threat hunts, which can improve an organization's abilities to counter future threats effectively.

The metrics reported in M-Trends 2025 are based on Mandiant Consulting investigations conducted between Jan. 1 and Dec. 31, 2024, that found targeted attack activity.

Campaigns & Global Events

Campaigns are a set of impactful intrusions conducted by an attacker or multiple attackers in cooperation toward a single objective at multiple targets within a relevant time frame.[1]

Global Events are a set of impactful intrusions conducted by multiple unrelated adversaries in parallel campaigns involving a similar theme, target, or resource.

When Mandiant experts identify threat activity that is actively impacting multiple organizations, a Campaign or a Global Event is created. Campaigns represent focused efforts by one or more threat groups with a single objective. Global Events encompass multiple threat groups pursuing different objectives but using similar tactics, such as exploiting a newly disclosed vulnerability. Mandiant delivers dynamic updates throughout the lifespan of each Campaign and Global Event, including details of indicators of compromise (IOCs) and tactics, techniques, and procedures (TTPs) unique to the event. Where possible, Mandiant provides examples, context, and information about threat actor behaviors, tools, and malware as well as actionable defensive and preventative measures. This intelligence is based on real-world data collected from Mandiant investigations and research, enabling our clients to respond effectively and decisively to active threats at first discovery and as they evolve.

In 2024, Mandiant initiated 83 campaigns and five global events and continued to track activity identified in previous years. These campaigns affected every industry vertical and 73 countries across six continents. Figure 1 depicts 33 campaigns and three global events, a subset of all campaigns and global events with direct relation to Mandiant incident response engagements.

For example, Campaign 23.042 began in April 2023 when the financially motivated group UNC3944 obtained network access to various organizations via SMS phishing and social engineering. With this access, UNC3944 ultimately stole proprietary data and deployed the ALPHV ransomware.

Other examples include Russian cyber espionage groups like APT28 and APT44. Campaign 23.056 tracked a subcluster of Russian cyber espionage group APT28 that, starting in late August, conducted credential harvesting and exploited Microsoft Outlook vulnerability CVE-2023-23397. Campaign 24.004 tracks APT44 activity leveraging trojanized software installers distributed via torrents on Ukrainian- and Russian-language forums as a means of achieving opportunistic initial access to potential targets of interest. In observed cases, victims of interest to APT44 received publicly available malware, such as DARKCRYSTALRAT,[2] for follow-on exploitation.

To facilitate tracking and analysis of large-scale events, such as widespread exploitation of a vulnerability, Mandiant utilizes global events as a framework to encapsulate multiple distinct campaigns. For instance, Global Event 24.004 groups three campaigns (CAMP.24.026, CAMP.24.030, CAMP.24.031) associated with different threat actors exploiting CVE-2024-3400. Each campaign tracks unique tactics, techniques, and procedures (TTPs), such as SNOWLIGHT downloader deployment, reconnaissance targeting configuration files, and BEACON backdoor usage. Global Event 24.002 tracks zero-day exploitation of CVE-2023-46805 and CVE-2024-21887, encompassing UNC5221 deploying custom malware and web shells, and another actor deploying SLIVER and TERRIBLETEA backdoors.

2024 Campaigns and Global Events Related to Mandiant Incident Response Investigations

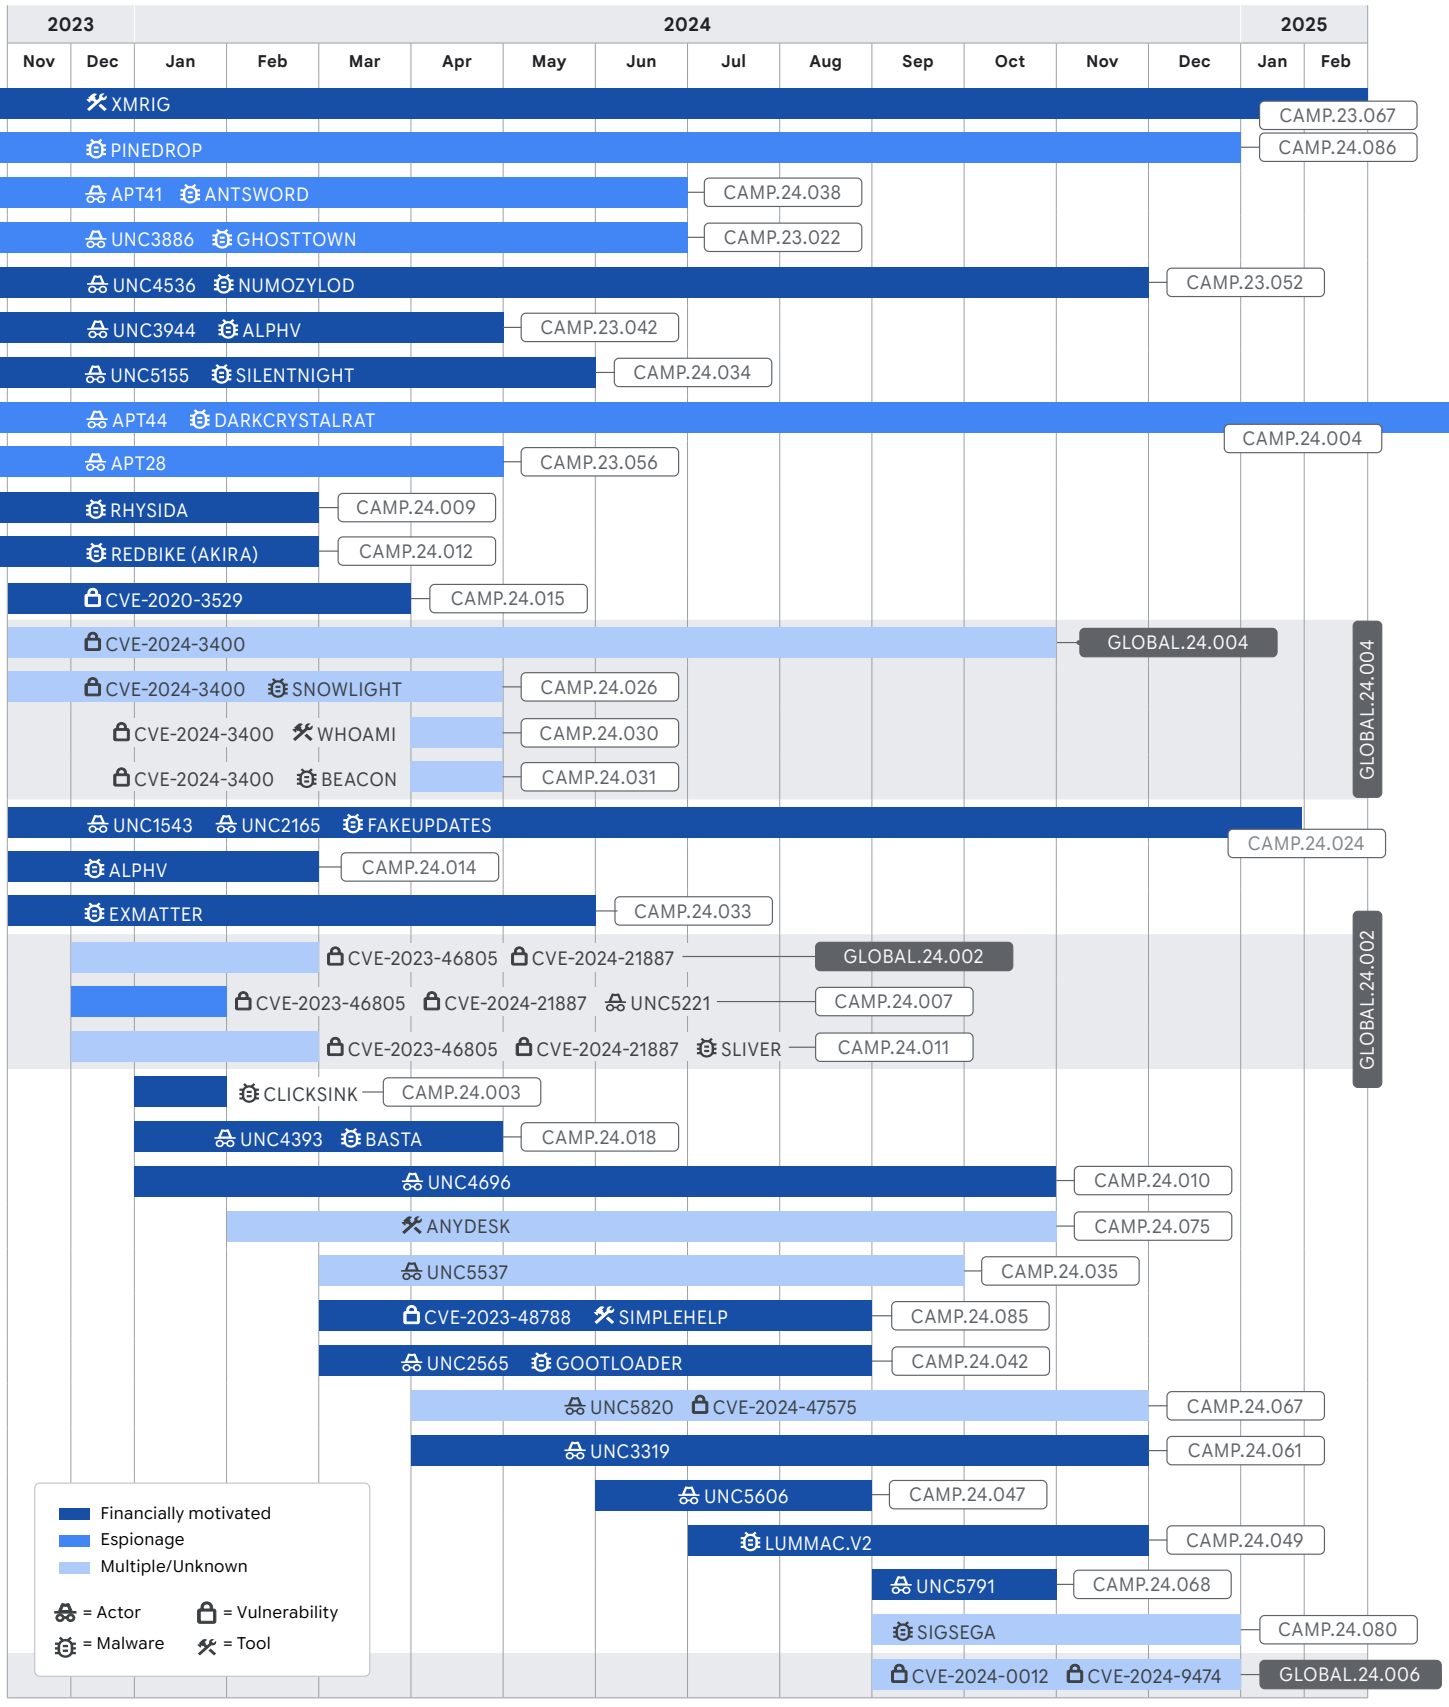

Figure 1: Campaigns and global events related to 2024 Mandiant incident response investigations

Targeted Attacks

Targeted Industries

An **industry category** describes an organization's primary industry. Organizations are typically assigned to only one category that best describes its primary industry, though many organizations have links to multiple industries. For example, a cryptocurrency exchange relates both to the financial and technology sectors, but for the purposes of this section, it would be categorized as a financial sector organization.

Mandiant responded to incidents affecting the financial sector more than any other sector in 2024. Business and professional services, high tech, government, and healthcare made up the next most frequently observed sectors. These top industries are consistent with prior years, with slight variations. For example, in 2023, investigations associated with retail and consumer goods and services organizations slightly outpaced those associated with healthcare and government entities, while the opposite was true in 2024.

Targeted Industries, 2024

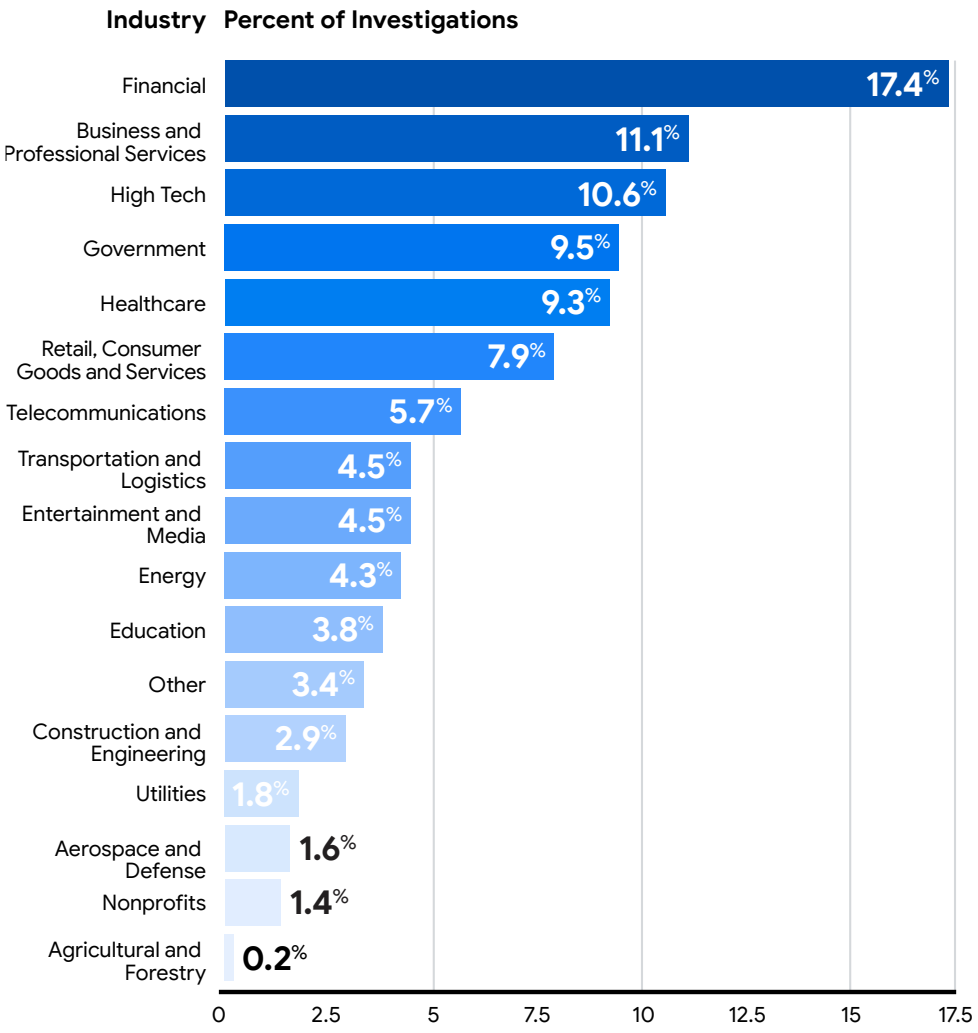

Initial Infection Vector

For the fifth year in a row, exploits were the most frequently observed initial infection vector in Mandiant incident response investigations. For intrusions in which an initial infection vector was identified, 33% began with exploitation of a vulnerability. This is a decline from 2023, during which exploits represented the initial intrusion vector for 38% of intrusions, but nearly identical to the share of exploits in 2022, 32%.

Initial Infection Vector, 2024

Stolen credentials overtook email phishing as the second most frequently observed initial infection vector in 2024, representing 16% of intrusions, compared to 14% for email phishing. In 2023, email phishing was determined to be the initial infection vector in 17% of intrusions and stolen credentials in just 10%. While email phishing remains a common and effective method for obtaining initial access, adversaries can obtain credentials in a variety of ways, including purchasing leaked or stolen credentials on underground forums, mining large data leaks for credentials, and actively pursuing credentials by infecting users with keyloggers and infostealers. The continued prevalence of phishing and credential theft underscores the importance of implementing multifactor authentication (MFA), preferably FIDO2-compliant MFA methods.

Phishing Declines as an Initial Infection Vector, 2022-2024

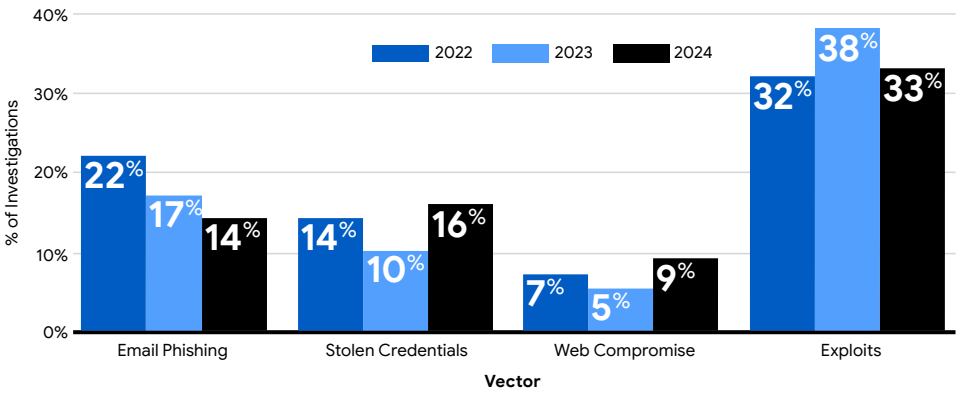

The percentage of intrusions that began with web compromise increased from 5% in 2023 to 9% in 2024. Web compromise encompasses drive-by compromise, the use of malicious advertisements, search engine optimization (SEO) poisoning, and compromised websites. To help mitigate risk from web compromise, organizations should consider a multilayered approach encompassing endpoint script blocking, content filtering for malicious redirects and software, policies against browser credential storage, and consistent patching of all systems.

In 2024, prior compromise remained a relatively common initial infection vector, occurring in 8% of investigations. The continued prevalence of this vector likely reflects the enduring effectiveness of threat actors specializing in establishing initial access, then providing that access to other threat actors.

Insider threat, typically a negligible proportion of Mandiant's incident response investigations, emerged as a surprisingly consequential initial infection vector in 2024. Specifically, a surge in North Korean IT workers seeking employment under false pretenses led to insider threat representing 5% of identified initial infection vectors. Mandiant primarily tracks this activity as UNC5267.

Mandiant also observed threat actors gain access to targeted systems through brute-force attacks, third-party compromise, social engineering voice calls (voice phishing or vishing), SIM swapping, supply chain compromise, and Bring Your Own Device (BYOD)—typically infected USBs.

Mandiant was unable to determine an initial infection vector for 34% of 2024 intrusions. Although numerous factors can contribute to an unknown vector, this considerable proportion indicates potential deficiencies in enterprise logging and detection capabilities.

Most Frequently Exploited Vulnerabilities

Among the Mandiant incident response investigations performed in 2024, the most frequently exploited vulnerabilities affected security devices, which are, due to their function, typically placed at the edge of the network. Three of the four vulnerabilities were first exploited as zero-days. While a broad selection of threat actors have recently targeted edge devices, Mandiant also specifically noted an increase[3] in targeting from Russian[4] and Chinese[5]cyber espionage actors.

Most Frequently Exploited Vulnerabilities

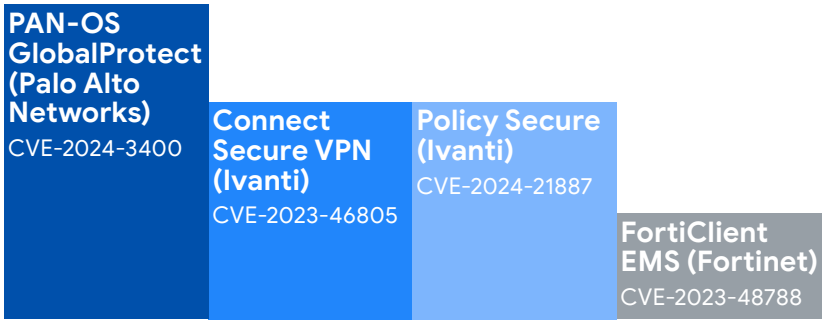

CVE-2024-3400

CVE-2024-3400 is a vulnerability in the GlobalProtect feature of Palo Alto Networks PAN-OS software that, when exploited, allows command injection through arbitrary file creation. Mandiant observed one threat group exploit this vulnerability as a zero-day. Within two weeks of its disclosure on April 12, 2024, and the publishing of proof-of-concept (PoC) code on April 13, 2024, Mandiant observed more than a dozen separately tracked groups exploiting this vulnerability, including a RANSOMHUB affiliate that used initial access established using this vulnerability to conduct multifaceted extortion.

CVE-2023-46805 and CVE-2024-21887

On Jan. 10, 2024, Ivanti disclosed[6] two vulnerabilities, CVE-2023-46805 and CVE-2024-21887, impacting Ivanti Connect Secure VPN ("CS," formerly Pulse Secure) and Ivanti Policy Secure appliances. Successful exploitation of these vulnerabilities allows authentication bypass and command injection, respectively. When chained together, these allowed for unauthenticated arbitrary command execution on systems. Mandiant identified[7] UNC5221, a suspected Chinese cyber espionage threat cluster, exploiting these vulnerabilities in the wild as zero-days as early as December 2023. UNC5221 leveraged multiple custom malware families, in several cases trojanizing legitimate CS files with malicious code. The malware functionality and observed activity suggest that UNC5221 was primarily focused on establishing persistent access, avoiding detection, and performing internal reconnaissance.

Ivanti worked closely with Mandiant, affected clients, government partners, and Volexity to address these vulnerabilities. They released a blog post with mitigations, patches, an enhanced external integrity checker tool,[8] and a disclosure for a subsequently discovered vulnerability, CVE-2024-21893. CVE-2024-21983 is a server-side request forgery vulnerability that allows a remote attacker to obtain unauthorized access. Mandiant also released a remediation and hardening guide.[9]

In mid-January 2024, Mandiant identified UNC5135 scanning Ivanti Connect Secure appliances but did not directly observe UNC5135 successfully exploit these vulnerabilities. Mandiant assesses with moderate confidence that UNC5135 is linked to UNC3236, which we suspect to align with the publicly reported Volt Typhoon.

By April 2024, Mandiant observed[10] eight distinct clusters involved in the exploitation of one or more of the three vulnerabilities: CVE-2023-46805, CVE-2024-21887, and CVE-2024-21893. Of these eight clusters, Mandiant tracked five suspected Chinese cyber espionage threat clusters that exhibited distinct post-compromise behavior and used different malware after exploiting the vulnerabilities for initial access.

CVE-2023-48788

CVE-2023-48788 is a SQL injection vulnerability in the FortiClient Endpoint Management Server. Mandiant observed a financially motivated threat cluster exploit this vulnerability to execute arbitrary SQL commands within two weeks of its March 12, 2024, disclosure. In observed operations, the threat cluster deployed the SimpleHelp remote administration tool, likely to establish persistent access before offering that access for sale to other threat actors.

In October and November 2024, a suspected FIN8 threat cluster gained access to a targeted organization by exploiting CVE-2023-48788, deployed SNAKEBITE ransomware, and used the publicly available backup utility RESTIC for data theft.

Global Detection by Source

Internal detection is when an organization independently discovers it has been compromised, such as through an internal security appliance alert or internal personnel notification of suspicious activity.

External notification is when an outside entity informs an organization it has been compromised, such as law enforcement agencies, cybersecurity companies, or industry partners (External Entity). In some cases, attackers will perform this notification, such as through a ransom note (Adversary).

The majority of organizations, 57%, first learned of a 2024 compromise from an external source. External notifications can be further divided into adversary notifications and external entity notifications. Adversary notifications typically take the form of ransom notes and represented 14% of total detection sources in 2024. Notifications from external entities, such as law enforcement or cybersecurity companies, comprised 43% of total detection sources. Organizations discovered an intrusion through internal mechanisms in 43% of 2024 investigations. These figures are roughly similar to our findings in 2023 investigations, which saw 54% external notifications and 46% internal notifications overall.

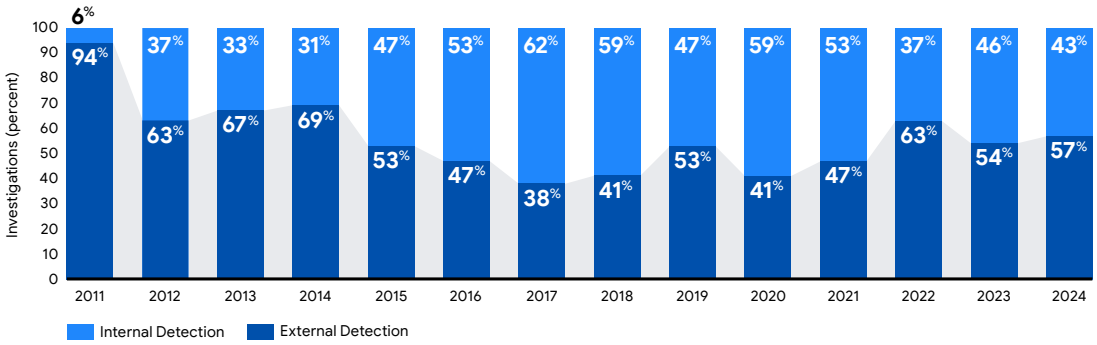

Global Detection by Source, 2024

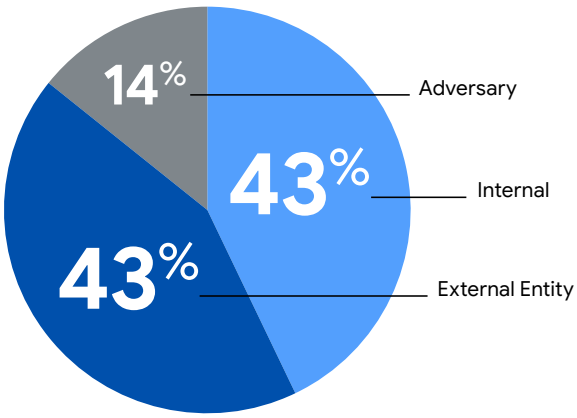

Global Median Dwell Time

Dwell time is calculated as the number of days an attacker is present in an environment that has been compromised before they are detected. The median represents a value at the midpoint of a data set sorted by magnitude.

The 2024 global median dwell time remained largely in line with 2023 figures. While the median overall value increased by one day from 2023 to 2024, the year-over-year trend continues to indicate that dwell times have declined significantly over the long term. For example, overall dwell time in 2014 was 205 days, compared to just 11 days in 2024. Dwell time for internally discovered intrusions remained less than that of all externally notified intrusions in 2024.

Median Dwell Time, 2011-2024

	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024
All	416	243	229	205	146	99	101	78	56	24	21	16	10	11
External	—	—	—	—	320	107	186	184	141	73	28	19	13	11
Internal	—	—	—	—	56	80	57.5	50.5	30	12	18	13	9	10

Median Dwell Time by Detection Source, 2024

	2024
All	11
Adversary	5
External Entity	26
Internal	10

The median adversary notification time was just five days, while external partners notified in a median of 26 days. This discrepancy is not surprising given that the vast majority of adversary notifications originate from extortion actors who benefit from monetizing intrusions quickly.

Global Dwell Time Distribution

Dwell time distribution plots intrusions that Mandiant investigated across ranges of dwell time. The distribution heat map demonstrates that the prevailing trend across Mandiant investigations from 2018 to 2024 is toward shorter and shorter dwell times. Comparing 2023 to 2024, the percentage of investigations that were discovered in one week or less increased from 43.3% to 45.1%.

Global Dwell Time Distribution, 2018-2024

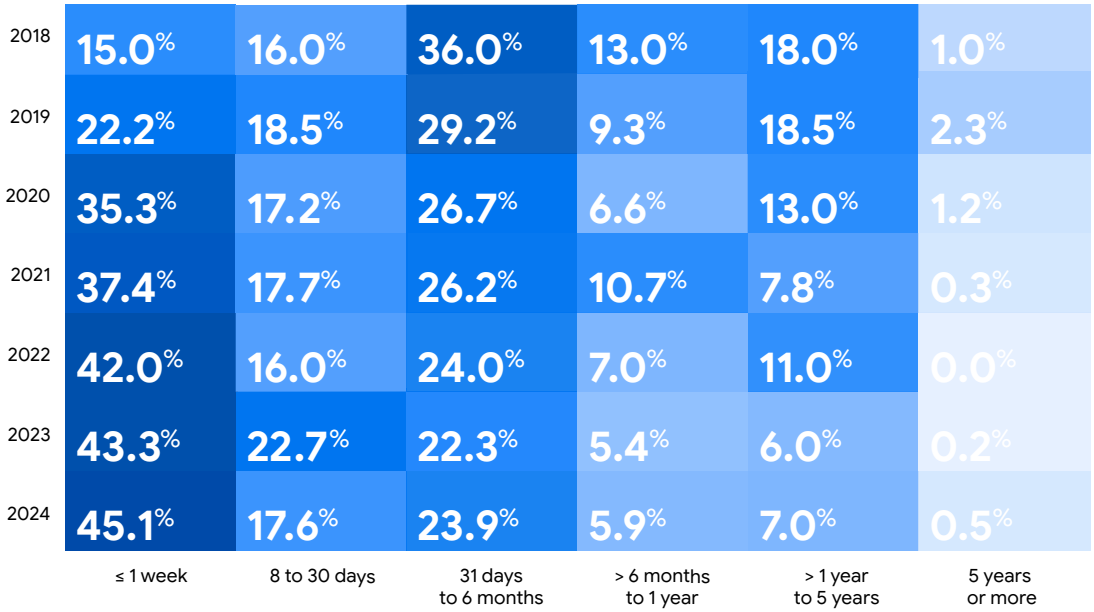

Post-Compromise Activity

Financial Gain, 2020-2024

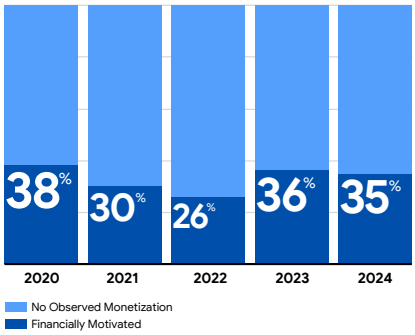

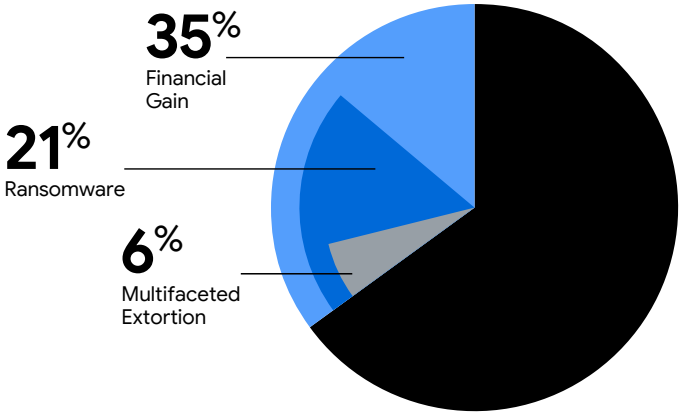

Financial Gain

In 2024, financially motivated intrusions, where a monetization technique was directly observed, represented 35% of all Mandiant incident response investigations. Ransomware-related intrusions represented 21% of all 2024 intrusions and approximately two-thirds of financially motivated intrusions. These proportions are also comparable to 2023, when ransomware was involved in 23% of all cases and about two-thirds of financially motivated intrusions.

In addition to ransomware-related events, Mandiant also responded to a variety of other financially motivated intrusions in 2024, including data theft extortion without ransomware encryption, illicit cryptomining, North Korean IT worker employment fraud, business email compromise, cryptocurrency theft, and cases in which threat actors attempted to monetize intrusions by offering access to targeted organizations or stolen data for sale.

Data Theft, 2020-2024

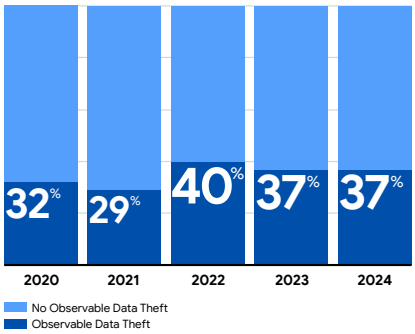

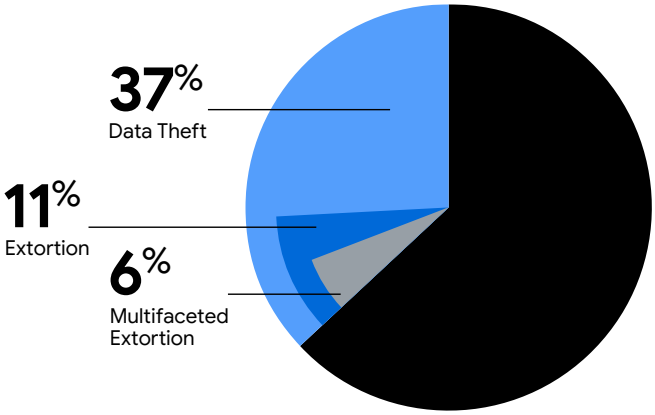

Data Theft

In 37% of 2024 investigations, Mandiant identified evidence of data theft, which is consistent with 2023. Data theft extortion events in which no ransomware was deployed represented 11% of all cases, and multi-faceted extortion, which includes both data theft and ransomware encryption, represents 6% of all cases.

Mandiant also observed attackers focus on theft of credentials and information useful for performing further reconnaissance of compromised networks. In addition, Mandiant identified attackers, such as the Russian cyber espionage actor APT28 and Chinese cyber espionage groups including APT41, conducting more targeted data theft. APT28 conducted selective data theft, demonstrating interest in personnel-related data, as well as email content and documents relevant to geopolitical topics consistent with Russian interests. In a campaign targeting multiple organizations in Europe, the Middle East, and Africa (EMEA) and Japan and Asia Pacific (JAPAC), APT41 leveraged SQLULDR2 to export data from Oracle Databases and used PINEGROVE to systematically and efficiently exfiltrate large volumes of sensitive data from the compromised networks, transferring to OneDrive to enable exfiltration and subsequent analysis.

Insider Threats

Mandiant responded to a number of incidents involving a unique variety of insider threat, North Korean IT workers. Mandiant primarily tracks this activity as UNC5267. North Korean IT workers use stolen and fabricated identities to apply for high-paying jobs in order to generate revenue for the North Korean regime in violation of international sanctions. Mandiant identified IT workers at diverse organizations, including in the financial services, telecommunications, media and entertainment, retail, and technology industries. In incident response engagements to date, North Korean IT workers have primarily functioned within the scope of their job responsibilities. However, the remote workers often gain elevated access to modify code and administer network systems. This heightened level of access granted to fraudulent employees presents a significant security risk. Moreover, in several cases in the latter half of 2024, Mandiant observed evidence of North Korean IT workers stealing proprietary data from targeted organizations and, following discovery and termination, threatening to release it publicly if the organization did not pay a ransom.

Mandiant released detailed guidance for detecting North Korean IT worker job applicants in Staying a Step Ahead: Mitigating the DPRK IT Worker Threat.[11]

Malware

A malware family is a program or set of associated programs with sufficient "code overlap" among the variants that Mandiant considers them to be largely the same thing, a "family." The term family broadens the scope of a single piece of malware as it can be altered over time, which in turn creates new, but fundamentally overlapping pieces of malware.

An **observed malware family** is a malware family identified during an investigation by Mandiant experts.

The **operating system effectiveness** of a malware is the operating system(s) that the malware can target.

In 2024, Mandiant began tracking 632 net new malware families. In investigations, Mandiant observed 205 malware families, 83 of which were both newly tracked and observed in at least one incident response investigation. This number of newly tracked families is on par with the 626 families Mandiant began tracking in 2023, bringing the total number of tracked malware families to more than 5,500 unique families. The 83 newly tracked families that Mandiant observed in incident response investigations in 2024 is lower than the 128 families observed in the same category in 2023. This continues a trend observed during the past three years of fewer new malware families being identified in investigations. This decrease showcases threat actors' continued willingness to leverage tools already present within the targeted environment as well as their ability to use and misuse tools rather than constructing new malware or configuring known post-exploitation tools. A growing number of compromises use no malware at all.

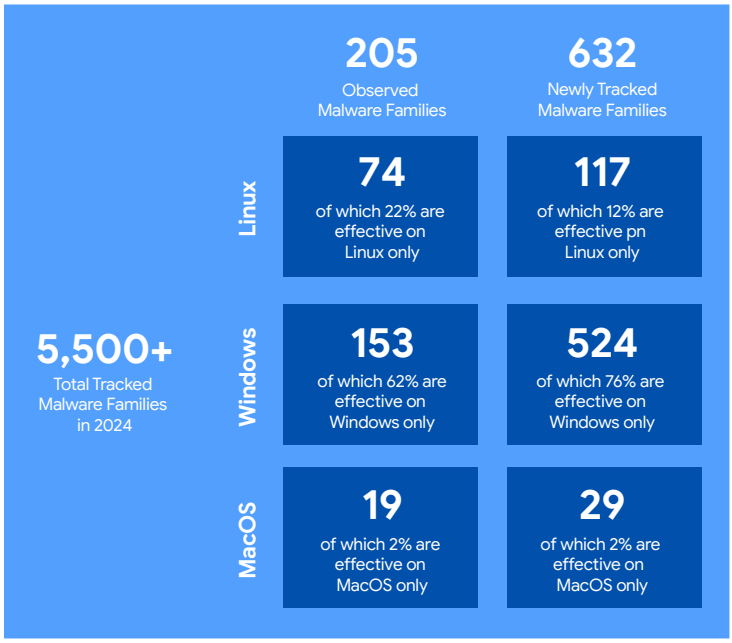

Looking further into the corpus of malware tracked by Mandiant, malware effective on Windows remains most prevalent. In both newly tracked (76%) and observed malware (62%) in 2024, Mandiant experts observed that malware was more likely to be effective exclusively on the Windows operating system. However, Mandiant has seen a decrease in the proportion of malware designed for Windows systems over the years.

Malware effective exclusively on Linux operating systems continues to increase slowly, accounting for 12% of newly tracked malware families and 22% of observed malware in 2024, compared to 11% of newly tracked and 17% of observed in 2023. The comparative reduction in Windows malware does not signify decreased risk associated with Windows systems but may indicate the risk to Linux environments is slowly increasing.

Malware Families by Category

A malware category describes a malware family's primary purpose. Each malware family is assigned only one category that best describes its primary purpose, regardless of functionality for more than one category.

Of the 632 malware families that Mandiant began to track in 2024, backdoors remained the predominant category, representing 31% of malware families. The next most observed categories were downloaders (19%), droppers (12%), credential stealers (6%), and ransomware (5%). The "Other" category is made up of utilities, tunnelers, data miners, rootkits, keyloggers, and point-of-sale malware, each of which make up less than 5% of the malware population. These findings continue to remain consistent year over year with little movement in position.

Newly Tracked Malware Families by Category, 2024

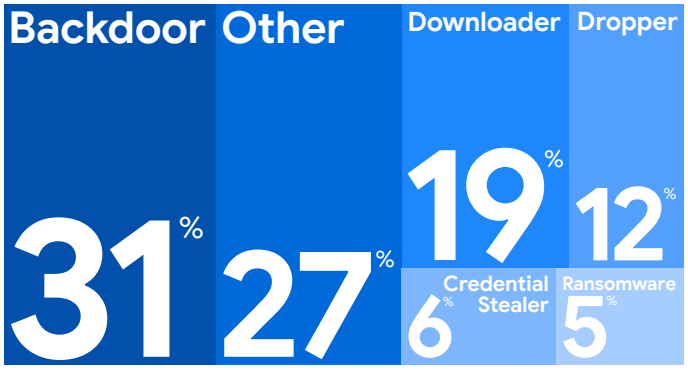

Similarly, observed malware family categories remained relatively consistent with the findings from previous years. Of the 205 unique malware families observed in investigations conducted during the 2024 calendar year, backdoors remained most used by attackers, with 35% of observed malware families with that primary purpose. The remaining malware family categories are made up of ransomware (14%), droppers (8%), downloaders (7%), tunnelers (6%), and credential stealers (5%).

In both the newly tracked and observed malware families by category, Mandiant continues to see a large portion of the percentage of malware residing in the "Other" category. This likely reflects the diversity of both attackers and objectives that Mandiant encounters in investigations.

Observed Malware Families by Category, 2024

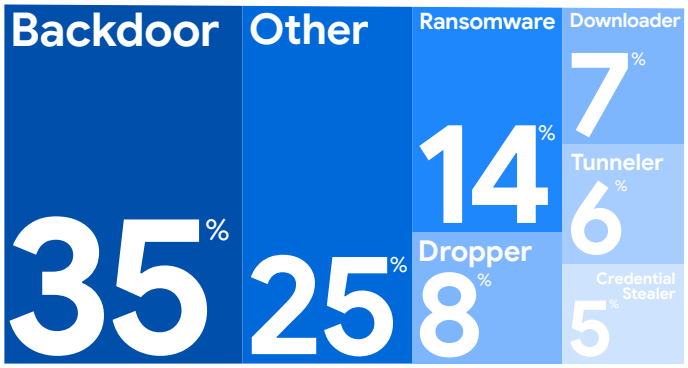

Malware Category	
Backdoor	A program whose primary purpose is to allow a threat actor to interactively issue commands to the system on which it is installed
Credential Stealer	A utility whose primary purpose is to access, copy, or steal authentication credentials
Data Miner	A utility whose primary purpose is to gather ("mine") data, typically for theft by threat actors. Excludes utilities that gather data such as credentials used for the purpose of escalating privileges or information used for system or network reconnaissance.
Downloader	A program whose sole purpose is to download (and perhaps launch) a file from a specified address, and which does not provide any additional functionality or support any other interactive commands
Dropper	A program whose primary purpose is to extract, install, and potentially launch or execute one or more files
Launcher	A utility with the primary purpose of gathering (or "mining") data, usually for theft by threat actors, excluding tools used solely for collecting privilege escalation credentials or reconnaissance information
Ransomware	A program whose primary purpose is to perform some malicious action (such as encrypting data), with the goal of extracting payment from the victim in order to avoid or undo the malicious action
Tunneler	A program that proxies or tunnels network traffic
Utility	A program that has a specialized purpose that does not fit into any other defined category (such as keylogger or sniffer)
Other	Includes all other malware categories such as utilities, tunnelers, data miners, rootkits, keyloggers, and point-of-sale malware

Most Frequently Seen Malware Families

For the fifth consecutive year, BEACON was identified as the most frequently observed malware family in Mandiant investigations globally and was identified in 5.4% of all intrusions. BEACON usage has decreased dramatically since 2021, when it was observed in 28% of Mandiant investigations.

Most Frequently Seen Malware Families, 2024

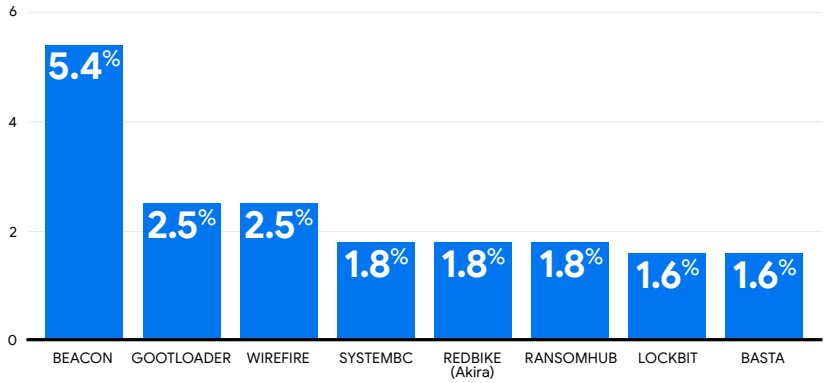

Of note, in July of 2024, Europol[12] provided an update on Operation MORPHEUS, a global action against the illicit use of the unlicensed versions of the Cobalt Strike red teaming tool. This operation, conducted with law enforcement and private sector partners, successfully disrupted infrastructure linked to cyber criminal activities. The initiative, which began in 2021, involved flagging 690 IP addresses, 593 of which were taken down by online service providers. Fortra,[13] the maintainers of the Cobalt Strike framework, also announced the number of unauthorized copies of Cobalt Strike observed in the wild has decreased by 80% over the past two years as a result of their participation in Operation MORPHEUS. Observed declines in percentages of investigations where Mandiant identified BEACON since 2021 may reflect the success of this effort.

BEACON Usage, 2020-2024

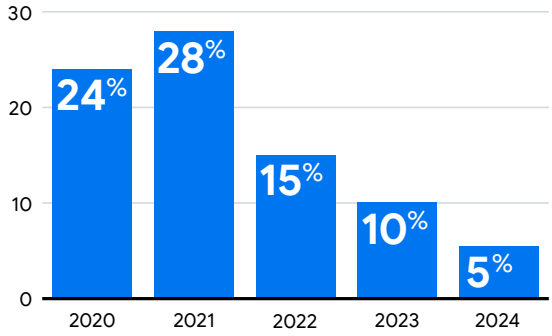

Malware Family	
BASTA	BASTA is a ransomware written in C++ that encrypts local files. The ransomware is capable of deleting volume shadow copies. BASTA generates a random ChaCha20 key to encrypt each file; the key is encrypted and appended to the end of the file. The malware has been observed using .basta as the extension for encrypted files; however, some samples have used a random nine-character alphanumeric extension.
BEACON	BEACON is a backdoor written in C/C++ that is part of the Cobalt Strike framework. Supported backdoor commands include shell command execution, file transfer, file execution, and file management. BEACON can also capture keystrokes and screenshots as well as act as a proxy server. BEACON may also be tasked with harvesting system credentials, port scanning, and enumerating systems on a network. BEACON communicates with a command-and-control (C2 or C&C) server via HTTP or DNS.
GOOT-LOADER	GOOTLOADER is a JavaScript downloader that comes in an obfuscated form. It downloads another JavaScript file that drops and executes the intended payload.
LOCKBIT	LOCKBIT is a ransomware written in C that encrypts files stored locally and on network shares. LOCKBIT can also identify additional systems on a network and propagate via SMB. Prior to encrypting files, LOCKBIT clears event logs, deletes volume shadow copies, and terminates processes and services that may impact its ability to encrypt files. LOCKBIT has been observed using the file extension ".lockbit" for encrypted files.
RANSOMHUB	RANSOMHUB is ransomware written in GoLang capable of encrypting data using ChaCha20, xChaCha20 or AES256 algorithms. The symmetric encryption key is per-file and protected by elliptic curve cryptography, ed25519. RANSOMHUB can be configured to encrypt a targeted directory, local disks, or network shares. RANSOMHUB provides the capability to reboot in safe mode before running or as a safe mode instance and can be configured for standard out logging.
REDBIKE	REDBIKE (also known as Akira) is ransomware written in C++ that encrypts local files. Encrypted files have the extension ".akira" appended to the filename. Files are encrypted using ChaCha20, and a ransom note is written to every folder with encrypted files. REDBIKE has some code overlaps with CONTI ransomware.
SYSTEMBC	SYSTEMBC is a tunneler written in C that retrieves proxy-related commands from a C2 server using a custom binary protocol over TCP. A C2 server directs SYSTEMBC to act as a proxy between the C2 server and a remote system. SYSTEMBC is also capable of retrieving additional payloads via HTTP. Some variants may utilize the Tor network for this purpose. Downloaded payloads may be written to disk or mapped directly into memory prior to execution. SYSTEMBC is often utilized to hide network traffic associated with other malware families. Observed families include DANABOT, SMOKELOADER, and URSNIF.
WIREFIRE	WIREFIRE is a web shell written in Python that exists as trojanized logic to a component of the Pulse Secure appliance. WIREFIRE supports downloading files to the compromised device and executing arbitrary commands.

Threat Groups

What is an UNC group?
When Mandiant encounters new threat activity that cannot confidently be linked to an existing group, an UNC group designation is created to tie together observable artifacts associated with the activity. As new information and artifacts are discovered that can be tied back to the same activity cluster, Mandiant analysts build on the initial understanding of the attacker, potentially merging it with other tracked threat clusters and ultimately graduating the UNC to an APT or FIN group.

In 2024, Mandiant identified and began tracking 737 new threat clusters, bringing the grand total of threat groups Mandiant tracks to more than 4,500. During 2024 incident response engagements, Mandiant observed 302 different threat groups, 233 of which were newly identified within the year. These figures are on par with 2023, during which Mandiant experts identified 719 new threat clusters and observed 316 groups in incident response investigations, with 220 of those groups also being newly identified.

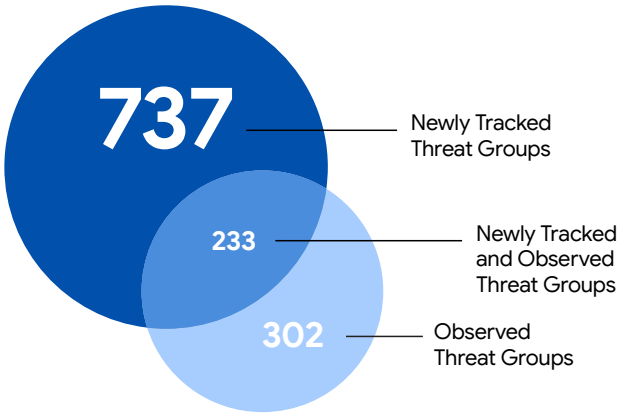

Organizations faced four advanced persistent threat (APT) groups from China, Russia, and Iran; one named financial threat (FIN) group; and 297 UNC groups from various geolocations in 2024 engagements. Mandiant continues to see groups that have been tracked for more than one year, and in some cases, up to 10 years. However, the majority of newly tracked and observed threat groups are new clusters of activity observed within Mandiant Consulting engagements in 2024. The composition of this set of threat clusters indicates that organizations continue to face a variety of both established and novel threats.

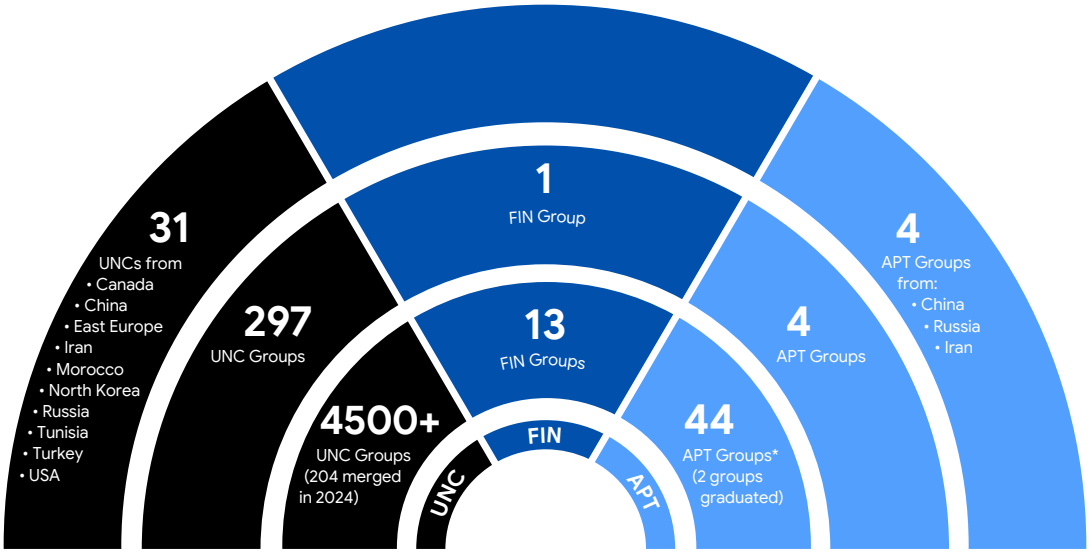

*Mandiant tracks Advanced Persistent Threat (APT) groups 0-45. Over the years, APT 11 and APT 13 were merged into other groups and subsequently deprecated resulting in 44 APT groups actively tracked by Mandiant.

Observed Groups by Goal

Observed Threat Groups by Goal, 2024

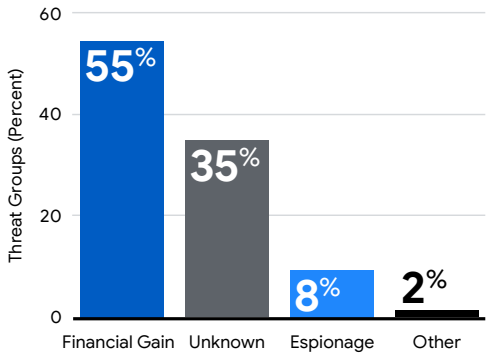

The majority of attackers active in 2024 were financially motivated (55%). This proportion is slightly larger than the 52% observed in 2023 and 48% observed in 2022. The growing share of financially motivated threat groups in Mandiant incident response investigations is likely due, in part, to the overall growth of impactful extortion intrusions. Espionage-motivated attackers represented 8% of threat groups identified in 2024 intrusions, compared to 10% in 2023. This is at least partially attributable to the number of distinct suspected Chinese cyber espionage activity clusters involved in vulnerability exploitation campaigns. A small percentage, 2%, included threat clusters Mandiant judged to be operating for hacktivist motivations and attackers focused on disruption or destruction. Several of these intrusions were linked to geopolitical motivations, including the conflicts in Ukraine and Gaza. Based on the evidence available at the time, Mandiant was unable to determine a motivation for the final 35% of groups.

Actor Graduations and Merges

In 2024, Mandiant graduated two new named threat groups, APT44 and APT45, and merged 204 activity clusters into other threat groups based on extensive research into activity overlaps. For details on how Mandiant defines and references UNC groups and merges, please see "How Mandiant Tracks Uncategorized Attackers."[14]

APT44

Sponsored by Russian military intelligence, APT44[15] (aka Sandworm, FROZENBARENTS) is a dynamic and operationally mature threat actor that is actively engaged in the full spectrum of espionage, attack, and influence operations. APT44 has aggressively pursued a multipronged effort to help the Russian military gain a wartime advantage and is responsible for nearly all of the disruptive and destructive operations against Ukraine over the past decade. APT44's support of the Kremlin's political objectives has resulted in some of the largest and most consequential cyberattacks in history. These operations include first-of-their-kind disruptions of Ukraine's energy grid in the winters of 2015 and 2016, the global NotPetya attack timed to coincide with Ukraine's Constitution Day in 2017, and the disruption of the opening ceremony of the 2018 Pyeongchang Olympics in response to Russia's doping ban from the games. Due to its history of aggressively using network attack capabilities across political and military contexts, APT44 presents a persistent, high-severity threat to governments and critical infrastructure operators globally where Russian national interests intersect.

APT45

Mandiant assesses with high confidence that APT45[16] is a moderately sophisticated cyber operator that supports the interests of the Democratic People's Republic of Korea (DPRK). Since at least 2009, APT45 has carried out a range of cyber operations aligned with the shifting geopolitical interests of the North Korean state. Although the group's earliest observed activities consisted of espionage campaigns against government agencies and defense industries, APT45 has expanded its remit to financially motivated operations, including targeting of the financial vertical; we also assess with moderate confidence that APT45 engaged in the development of ransomware. In 2019, APT45 directly targeted nuclear research facilities and nuclear power plants, such as the Kudankulam Nuclear Power Plant in India, marking one of the few publicly known instances of North Korean cyber operations targeting critical infrastructure.

Ransomware

A ransomware-related intrusion provides access for or is associated with a malicious actor that has the primary goal of encrypting data with the intention of extracting payment from the target.

Ransomware, data theft extortion, and multifaceted extortion are and will continue to be the most disruptive type of cyber crime globally, both due to the volume of intrusions and the scope of potential damage for each event. The impact of ransomware and extortion operations extends far beyond the initial victim. Mandiant responded to ransomware-related intrusions affecting healthcare, local government, energy, high tech, education, financial sector organizations, and others across JAPAC, EMEA, and the Americas. Ransomware-related events accounted for just over one-fifth (21%) of all Mandiant incident response investigations in 2024.

Initial Infection Vector

In contrast to the overall dataset, the most commonly observed initial infection vector for ransomware-related intrusions, when the vector could be identified, was brute-force attacks. Password spraying, virtual private network (VPN) devices compromised through default credentials, and high-volume Remote Desktop Protocol (RDP) login attempts are examples of the types of brute-force attacks that Mandiant observed in 2024. Use of this tactic reinforces the importance of auditing and configuring internet-exposed infrastructure to require multifactor authentication (MFA), to require verification for remote attempts to register MFA on an account for the first time, and to lock accounts after a certain number of failed login attempts.

Stolen credentials and exploits were tied for the second most common initial infection vector for 2024 ransomware-related intrusions at 21% each, followed by prior compromise at 15%, and third-party compromise at 10%.

Initial Infection Vector, 2024
Ransomware-Related

Detection by Source

Detection by external sources was more common for ransomware-related than non-ransomware related intrusions, with notifications directly from adversaries representing the majority of the variance. This is consistent with the extortion business model in which attackers intentionally and abruptly notify organizations of a ransomware intrusion and demand payment. In 2024, adversaries notified organizations of ransomware-related compromises in 49% of cases, other external entities in 21% of cases, and organizations discovered compromises internally in 30% of cases. In investigations without a ransomware component, adversaries represented only 5% of detection sources, while other external entities notified in 48% of cases, and organizations identified evidence of malicious behavior for themselves in 47% of cases.

Detection by Source, 2024

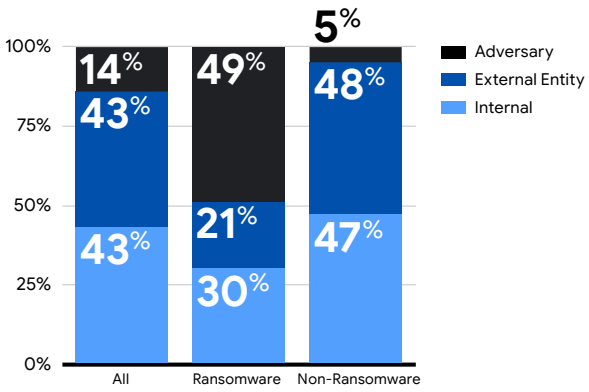

These figures are largely consistent with Mandiant's 2023 findings. External notifications in 2023 were also more common for ransomware-related intrusions (70%) than non-ransomware related intrusions (50%). Adversary notifications in 2023 represented approximately three quarters of external notifications for ransomware-related intrusions, while in 2024, the proportion of adversary notifications declined slightly to seven out of 10 of all external notifications for ransomware-related events.

Ransomware-Related Dwell Time vs Global Dwell Time

Median dwell time for ransomware-related intrusions was 11 days overall, five days for adversary-notified events, five days for compromises discovered by external entities such as law enforcement and cybersecurity companies, and 29 days for intrusions discovered internally.

Dwell Time Distribution for Ransomware-Related Intrusions

The dwell time distribution for ransomware-related intrusions is even more concentrated toward shorter time intervals between the first evidence of malicious activity and discovery of the incident. Events with a week or less of dwell time represent 56.5% of the ransomware-related intrusions that Mandiant investigated in 2024, compared to 45.1% of all intrusions discovered within one week. This finding is consistent with the extortion business model, in which attackers are incentivized to complete their objectives without being detected and swiftly and abruptly call the target organizations' attention to their activities.

Global Dwell Time Distribution, 2024

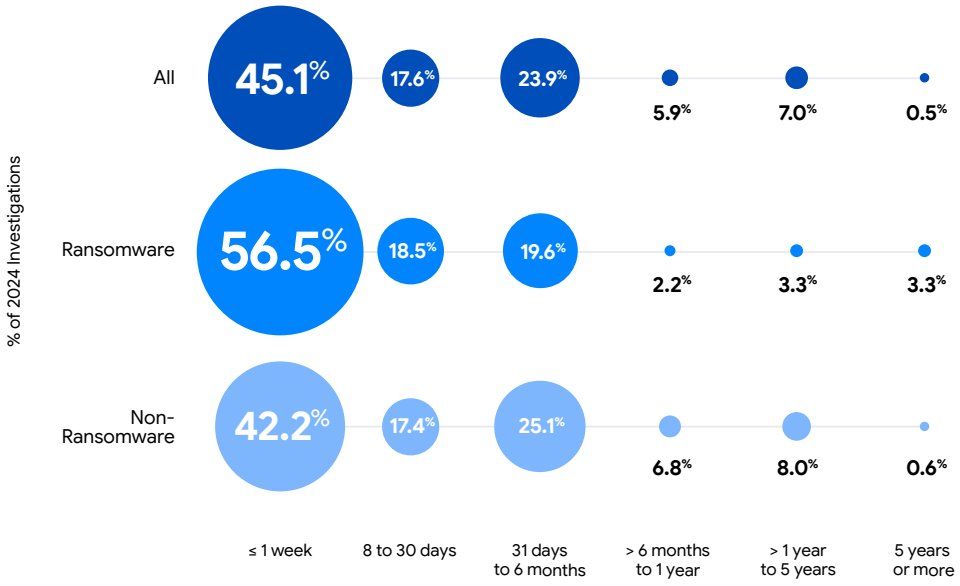

Malware

Unsurprisingly, the top malware category observed in 2024 ransomware intrusions was ransomware, which made up 34% of the malware data set. The next most prevalent categories are in line with the overall malware landscape observed in 2024. Credential stealers made up 12% of malware observed in ransomware-related intrusions, followed by backdoors (10%), utilities (10%), tunnelers (7%), downloaders (6%), and droppers (5%). The other 16% of malware families had other primary purposes such as keyloggers, launchers, installers, and uploaders.

Observed Malware Families by Category, 2024
Ransomware-Related

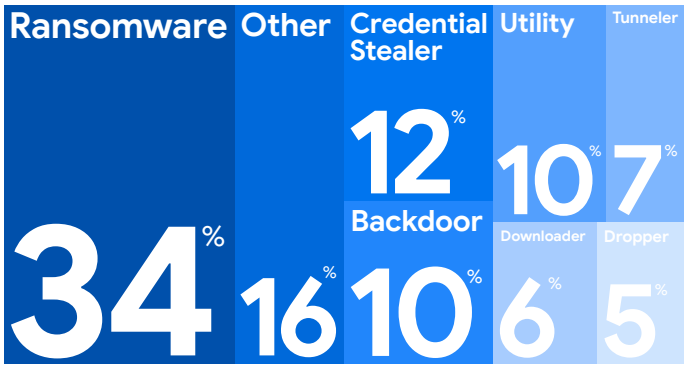

Compared to the overall metrics, ransomware-related intrusions saw a higher percentage of BEACON usage (15%). However, that may be attributable to the bias of the smaller dataset of ransomware-related intrusions rather than a true increase in the rate of BEACON usage when compared to all investigations. The next four most frequently observed malware families were ransomware varieties: RANSOMHUB (10%), REDBIKE (aka Akira) (10%), BASTA (9%), and LOCKBIT (9%). The tunneler SYSTEMBC (7%) was the sixth most commonly observed malware in ransomware-related intrusions, though it was the fourth most commonly observed family in all investigations. Several of these also appear in the overall most frequently seen malware families: BEACON, RANSOMHUB, REDBIKE, BASTA, LOCKBIT, and SYSTEMBC. The overlap of most frequently seen families for both overall and ransomware-related intrusions highlights how pervasive and prolific ransomware-related intrusions are.

Most Frequently Seen Malware, 2024
Ransomware-Related

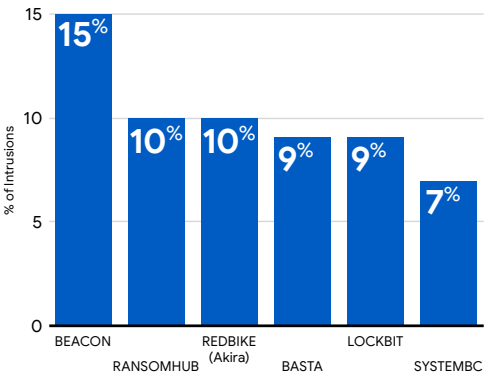

Compared to global metrics, ransomware-related intrusions saw more malware category variation; however, the ransomware-related malware dataset contains a much smaller proportion of backdoors. Within ransomware operations, this likely coincides with threat actors continuing to rely on remote control and administration tools. Credential stealers also make up double the percentage of the dataset in ransomware-related intrusions in 2024 compared to the overall dataset. Threat actors using ransomware are more likely to rely on publicly available and legitimate tools, such as credential extraction tools (credential stealers and remote administration tools), to accomplish their objectives.

Threat actors that conduct ransomware-related intrusions often rely on commercially available or legitimate tools to facilitate operations. This affords threat actors with various opportunities to blend in with the target environment, presumably delaying detection and therefore leading to more successful ransomware deployments against targets. Of these commercially available or legitimate tools, Mandiant observed that 37% of the tools used during intrusions in 2024 were utilities. This category includes utilities such as PsExec. Credential stealers made up nearly a fifth (18%) of tools observed in 2024 intrusions. Remote control and administration tools captured 12% of tools observed, followed by reconnaissance tools (11%) and tunnelers (8%). The remaining 14% of tools observed in ransomware-related intrusions fall into categories such as cryptomining tools, data mining tools, or tools used for lateral movement.

Observed Tools, 2024
Ransomware-Related

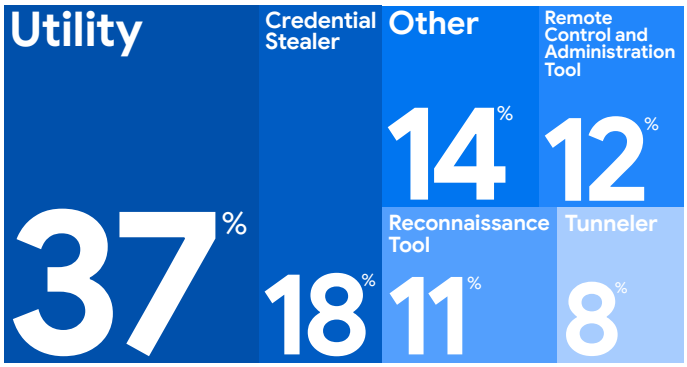

Tool Category	
Credential Stealer	A utility whose primary purpose is to access, copy, or steal authentication credentials
Reconnaissance Tool	A program whose primary purpose is to conduct some type of system or network reconnaissance (for example, enumerating accounts or systems, or conducting port scanning)
Remote Control and Administration Tool	A legitimate program whose primary purpose is to remotely access and control or administer a system
Tunneler	A program that proxies or tunnels network traffic
Utility	A program that has a specialized purpose that does not fit into any other defined category (such as keylogger or sniffer)
Other	Includes all other tool categories such as cryptomining tools, data mining tools, or tools used for lateral movement

Tools observed in 2024 ransomware intrusions were most frequently designed for the Windows operating system, almost certainly due to the operating system's high market share on desktops. SoftPerfect Network Scanner (NETSCAN), a network administration tool for Windows, mac OS, and Linux as well as PSEXEC, a Windows-native utility used to execute processes and launch interactive command prompts on other systems, were both observed in 29% of intrusions. NLTEST (19%) is often leveraged in ransomware deployment scripts or used manually by threat actors in the internal reconnaissance stage of the Targeted Attack Lifecycle,[17] as it is designed to help system administrators maintain domain controllers and active directory domains services, which serve as a main target in ransomware-related intrusions. The remainder of these tools are also publicly available—MIMIKATZ (18%), RCLONE (17%), ADVIPSCAN (15%), and AnyDesk (15%).

Most Frequently Seen Tools, 2024
Ransomware-Related

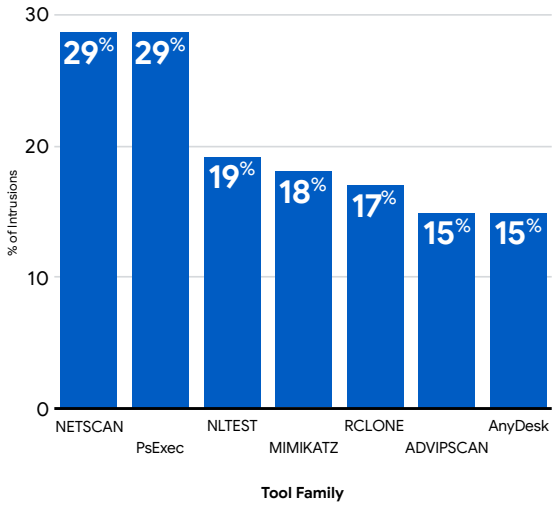

Tools

ADVIPSCAN	ADVIPSCAN is a publicly available network scanner developed by Famatech that has remote control capabilities.
AnyDesk	AnyDesk is a commercially available remote monitoring and management (RMM) application that is supported on Windows, macOS, Linux, Android, and ChromeOS devices.
MIMIKATZ	MIMIKATZ is a credential stealer written in C that targets Windows authentication credentials. Techniques employed include stealing password hashes, keys, and Kerberos tickets. Credentials can be printed to the console or saved to disk. MIMIKATZ also supports privilege escalation, extracting credentials from the Windows Local Security Authority Subsystem Service (LSASS) and Security Account Managers (SAM) database, and service manipulation.
NETSCAN	NETSCAN, the SoftPerfect Network Scanner, is a free multi-threaded IPv4/IPv6 scanner that pings computers, scans for listening TCP/UDP ports, discovers shared folders, and retrieves information about network computers via WMI, SNMP, HTTP, and NetBios.
NLTEST	NLTEST is the Microsoft nltest.exe utility, a command-line tool that is built into Windows Server 2008 and Windows Server 2008 R2.
PsExec	The PsExec utility, developed by Mark Russinovich as part of Sysinternals, is available from Microsoft.
RCLONE	RCLONE is a publicly available command-line utility to sync files and directories to and from numerous cloud-based resources, such as Amazon Drive, Dropbox, FTP, Google Drive, HTTP, Mega, Microsoft OneDrive, rsync.net, SFTP, and the local file system.

Ransomware Operations

RANSOMHUB

Data leak sites (DLS) are websites that publish stolen data of companies that refuse to pay a ransom. While this data is skewed toward targets who refused to pay attackers' ransom demands, it is still useful for understanding broad trends in extortion operations.

The RANSOMHUB ransomware-as-a-service (RaaS) and associated DLS launched in early 2024. By the second half of 2024, RANSOMHUB RaaS became the most prolific DLS that Mandiant tracks, taking the top spot from LockBit after its activity declined following law enforcement action. RANSOMHUB was also tied for most frequently observed ransomware in Mandiant incident response investigations performed in 2024. Mandiant currently tracks multiple threat clusters that have used this ransomware brand, including UNC2165, UNC5227, and others.

LockBit vs. RANSOMHUB DLS Listings, 2024

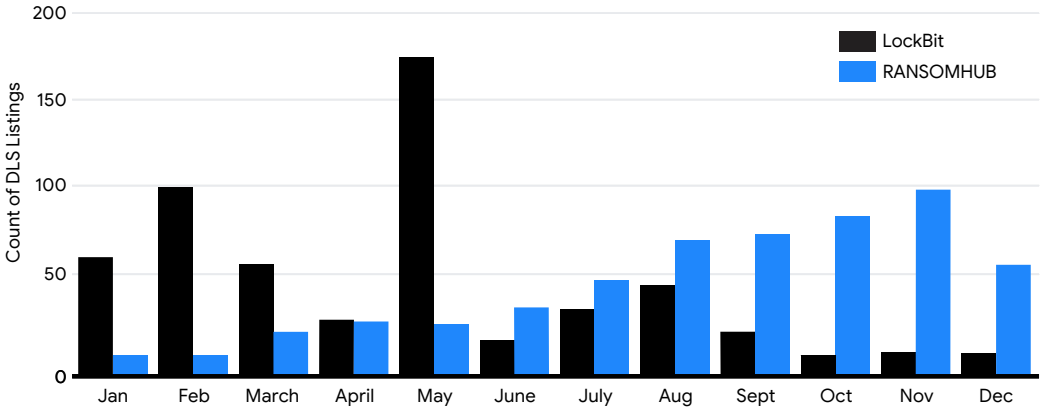

UNC2165[18] is a financially motivated threat cluster that has been active since at least 2019 and has conducted ransomware and data theft extortion operations using HADES, LOCKBIT, CONTI, and RANSOMHUB ransomware. UNC2165 has primarily gained access to victim organizations from FAKEUPDATES infections, although, since late 2020, some intrusions appeared to leverage stolen credentials. UNC2165 has used various methods to escalate privileges conducting Mimikatz and Kerberoasting attacks, targeting authentication data stored in the Windows registry, and searching for documents or files associated with password managers or that may contain plaintext credentials. Historically, UNC2165 operations heavily relied on BEACON for lateral movement and to maintain access to the victim environment; however, since late 2023, UNC2165 has used the MYTHIC post-exploitation framework and VIPERTUNNEL tunneler in intrusions. In most cases, UNC2165 has also stolen data from victims using Rclone or MEGASync.

UNC5227 is a financially motivated threat cluster active since at least November 2023 that has monetized access via ransomware deployment and data theft extortion. In some cases, UNC5227 has gained access to victim networks via brute-force attacks or stolen VPN credentials obtained from a separate threat cluster. UNC5227 relies on open-source tools, including MIMIKATZ and OPENSSH, to compromise additional accounts and move laterally through the network. They have also used PORTLIGHT, a custom Windows PowerShell utility for port-forwarding access using SecureShell (SSH) to maintain persistence, which may be exclusive to UNC5227. UNC5227 also uses EXMATTER, a private file upload tool, on compromised devices for data staging and theft before deploying ransomware. UNC5227 has deployed LOCKBIT.BLACK, ALPHV, RHYSIDA, and RANSOMHUB ransomware, based on direct observations as well as overlaps observed in the wild with EXMATTER and reverse Secure Shell (SSH) infrastructure.

REDBIKE (aka Akira)

The REDBIKE (aka Akira) RaaS first emerged in early 2023 and has remained one of the most active based on the quantity of successfully compromised organizations posted to its DLS. REDBIKE matched RANSOMHUB for most frequently observed ransomware in Mandiant incident response investigations performed in 2024. Mandiant tracks multiple threat clusters that have deployed this ransomware, including UNC5277 and UNC5280.

UNC5277 is a financially motivated threat cluster that has deployed REDBIKE ransomware in extortion operations involving both Windows and ESXi environments. In intrusions where the initial access vector is known, UNC5277 has leveraged stolen credentials to gain access to victim VPNs and has relied on publicly available tools to perform internal reconnaissance, escalate privileges, and maintain a presence in the environment. UNC5277 has used FORGEDGRIT, a public exploit for CVE-2023-27532, to steal credentials from Veeam backup servers in multiple intrusions. This threat cluster has stolen data via WinSCP for use in data theft extortion attempts.

UNC5280 is a financially motivated threat cluster active since at least December 2023 that has deployed REDBIKE ransomware and engaged in data theft operations. UNC5280 has leveraged valid VPN credentials to gain access to victim environments. UNC5280 initiated a SSH connection via FreeSSHd or MobaXterm and likely transferred REDBIKE samples to other hosts. Prior to the deployment of REDBIKE, UNC5280 has used Metasploit and surveyed target systems to exfiltrate both data and credentials. The threat cluster has also deleted forensic artifacts.

Cloud Compromises

Cloud compromises consist of intrusions where threat actors access a target's cloud environment, excluding the misuse of cloud services for attacker operations or infrastructure such as staging payloads or data theft.

In 2024 investigations, Mandiant observed threat actors compromise cloud assets through a variety of means. The most commonly observed initial infection vectors included email phishing (39%), stolen credentials (35%), SIM swapping (6%), and voice phishing or vishing (6%). Mandiant also noted use of prior compromise, exploits, third-party compromise, brute-force attacks, and malicious insiders—specifically North Korean IT workers applying for jobs under false pretenses—in order to gain access to cloud systems.

Cloud Initial Infection Vectors, 2024

In terms of objectives, data theft was observed in nearly two-thirds of cloud compromises (66%). Over a third of cases (38%), served financially motivated goals, including data theft extortion without ransomware encryption (16%), business email compromise (BEC) (13%), ransomware (9%), as well as cryptocurrency theft and employment fraud.

Two of the most frequently observed threat actors in cloud intrusions were UNC3944 and UNC5537.[19] Beginning in spring 2024, UNC5537[20] used stolen credentials to gain access to data belonging to clients of the Snowflake cloud data warehousing platform. The threat actor downloaded data and attempted to extort targeted organizations or sell the data on cyber crime forums. Mandiant found no evidence that a breach of Snowflake's environment occurred, only Snowflake client credentials.

UNC3944[21] used persistent social engineering techniques to gain access to targeted organizations, often calling service desks and convincing staff to reset passwords and multi-factor authentication (MFA) methods, including for privileged accounts. After obtaining access, Mandiant observed UNC3944 use a number of techniques to manipulate cloud hosted systems and services. The threat actor abused single sign on (SSO) solutions, for example assigning a compromised account to every application linked to an SSO instance, expanding the scope of the intrusion beyond on-premises infrastructure to cloud and SaaS applications. Mandiant identified UNC3944 using SSO applications to create new virtual machines (VMs), which they used to conduct follow-on activities. UNC3944 used compromised accounts to identify and access a variety of additional SaaS applications. In at least one case, UNC3944 used RANSOMHUB ransomware to encrypt an organization's virtualized environment. UNC3944 also abused cloud synchronization utilities, to move data from cloud-hosted data sources in the targeted environment to external attacker-owned cloud storage resources.

Threat Techniques

MITRE ATT&CK® is a globally accessible knowledge base of adversary tactics and techniques based on real-world observations. The ATT&CK knowledge base is used as a foundation for the development of specific threat models and methodologies in the private sector, government, and the cybersecurity product and service community.

Since the M-Trends 2020 report, Mandiant has supported the security industry by aligning its findings with the MITRE ATT&CK framework. To help organizations bolster their security, Mandiant provides metrics around the most commonly observed adversary tactics and sub-techniques. This information can enable organizations to prioritize the development of detection capabilities that address these prevalent threats, then inform strategic decisions on further security planning to improve security capabilities.

In October 2024, MITRE released ATT&CK framework version 16.1, which aligned techniques and sub-techniques to better reflect real-world adversary activity and improved platform descriptions. This change did not introduce a significant number of new techniques and sub-techniques to the already established framework. Mandiant began tracking two new ATT&CK techniques and 29 new sub-techniques in 2024 and mapped an additional 570 Mandiant techniques to the MITRE ATT&CK framework. Mandiant now tracks over 4,000 Mandiant Techniques that map to the ATT&CK framework, which totals 203 techniques and 456 sub-techniques. The observed MITRE ATT&CK techniques mapped to the Mandiant Targeted Attack Lifecycle can be found in the appendix of this report.

MITRE ATT&CK Techniques Used Most Frequently

Mandiant experts observed adversaries use 71% of MITRE ATT&CK techniques and 40% of sub-techniques during 2024 intrusions. This is relatively consistent with the two previous M-Trends reporting periods, during which nearly three-fourths of techniques and nearly half of sub-techniques were actively observed by Mandiant experts.

MITRE ATT&CK techniques in 2024 largely mirrored those of 2023, showing that these techniques have remained remarkably stable for several years. In nearly half of investigations, Mandiant investigators noted the use of a command or scripting interpreter (T1059) by attackers. Notable divergences from 2023 relate to Data Encrypted for Impact (T1486) and the use of External Remote Services (T1133). Data Encrypted for Impact (T1486) appears for the first time in the top 10 most frequently used techniques in 2024, indicating continued popularity of ransomware operations. While the use of Remote Services (T1027) has remained in the top 10 techniques for the past three years, the notable differences between Remote Services (T1027) and External Remote Services (T1133) lie within their definitions. Remote Services (T1027) relates to an attacker moving laterally through an environment with valid credentials, using system-based services that accept remote connections, which has been a typical attacker tactic over the years. The use of External Remote Services (T1133), or an adversary leveraging external-facing remote services such as virtual private networks (VPNs), Citrix, or other mechanisms to gain initial access to an environment, has been a focus for a number of threat clusters that Mandiant has tracked for years. However, it became popular among threat actors deploying ransomware throughout 2023 and 2024 and is now reflected in the M-Trends dataset.

Top 10 Most Frequently Seen MITRE ATT&CK Techniques

Rank	Technique	Percent
1	T1059: Command and Scripting Interpreter	44.6%
2	T1027: Obfuscated Files or Information	37.3%
3	T1021: Remote Services	35.3%
4	T1083: File and Directory Discovery	34.2%
5	T1070: Indicator Removal	29.4%
6	T1082: System Information Discovery	26.0%
7	T1140: Deobfuscate/Decode Files or Information	24.7%
8	T1486: Data Encrypted for Impact	22.9%
9	T1071: Application Layer Protocol	22.4%
9	T1133: External Remote Services	22.4%

Top 5 Most Frequently Seen MITRE ATT&CK Sub-Techniques

Rank	Technique	Percent
1	T1059.001: PowerShell	26.2%
2	T1021.002: SMB/Windows Admin Share	23.3%
3	T1021.001: Remote Desktop Protocol	22.6%
4	T1070.004: File Deletion	21.7%
5	T1569.002: Service Execution	19.0%

Regional Reports

Americas

The metrics reported in this section are based on Mandiant Consulting investigations affecting organizations that are located in North, Central, or South America.

Targeted Attacks

Initial Infection Vector

For compromises in the Americas in 2024 in which Mandiant was able to determine an initial infection vector, the most commonly observed vectors were exploits (28%), followed by stolen credentials (18%) and email phishing (16%). The distribution of initial infection vectors for the Americas is similar to what Mandiant observed globally in 2024 investigations.

AMERICAS

Exploit
28%

Stolen Credentials
18%

Email Phishing
16%

Detection by Source

In 2024 Mandiant investigations in the Americas, organizations were first notified of malicious activity in their environments by external parties 54% of the time and discovered evidence of suspicious activity internally 46% of the time. External notifications can be divided into 36% coming from external partners such as law enforcement and cybersecurity companies and 18% coming from attackers, largely in the form of ransom notes. These proportions are largely consistent with global figures for 2024.

Americas Detection by Source, 2017-2024

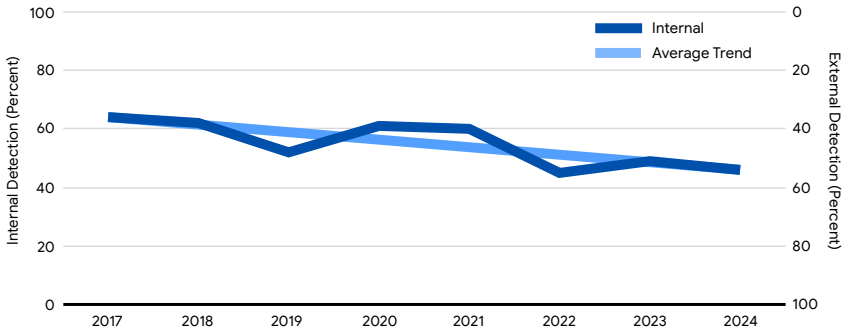

For 2024 ransomware-related intrusions in the Americas, adversaries first notified organizations of a compromise in 62% of cases, while external partners such as law enforcement or cybersecurity companies informed organizations in 9% of cases. Organizations discovered evidence of a ransomware-related incident internally in 29% of cases. This frequent rate of adversary notifications reflects the nature of extortion operations, which require contacting impacted organizations to initiate ransom negotiations.

Compared to global ransomware-related intrusion numbers, the Americas experienced higher rates of adversary notifications (62% compared to 49%) and lower rates of external partner notifications (9% compared to 21%). It is possible that the quantity of ransomware and extortion operations in North America accounts for this difference—the high volume of adversary activity is great enough that adversary notifications outpace external entity notifications by a larger margin in the Americas than globally. According to extortion data leak site (DLS) listings, the United States and Canada represent the first and third largest share of organizations, with United States organizations alone comprising half of all DLS listings.

Americas Detection by Source, 2024

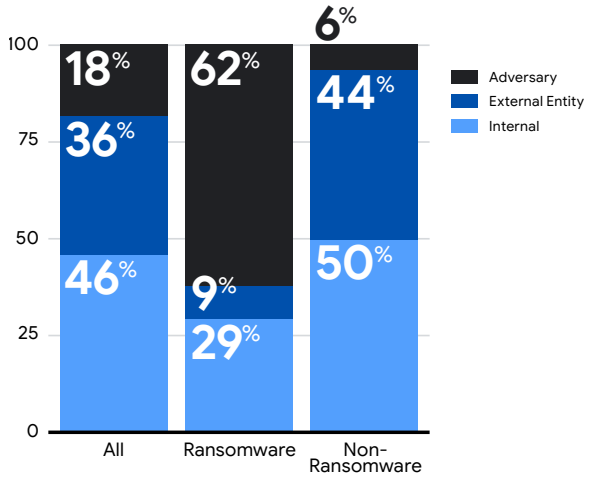

DLS Listings for the US vs. All Other Countries, 2020-2024

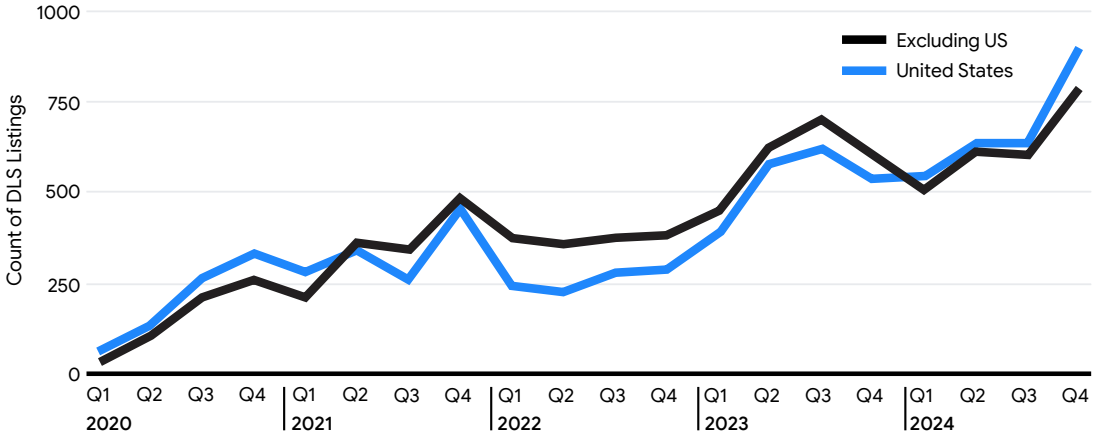

Median Dwell Time

The median dwell time for intrusions Mandiant investigated in the Americas in 2024 was 10 days overall, matching the median dwell time for 2023 and 2022. The median dwell time for internally and externally notified events in 2024 was also 10 days, which is also fairly consistent with prior years' data from the Americas as well as global trends. For ransomware-related events in the Americas in 2024, the median dwell time was six days versus 12 days for non-ransomware-related events. These numbers are similar to global numbers.

Americas Median Dwell Time, 2016-2024

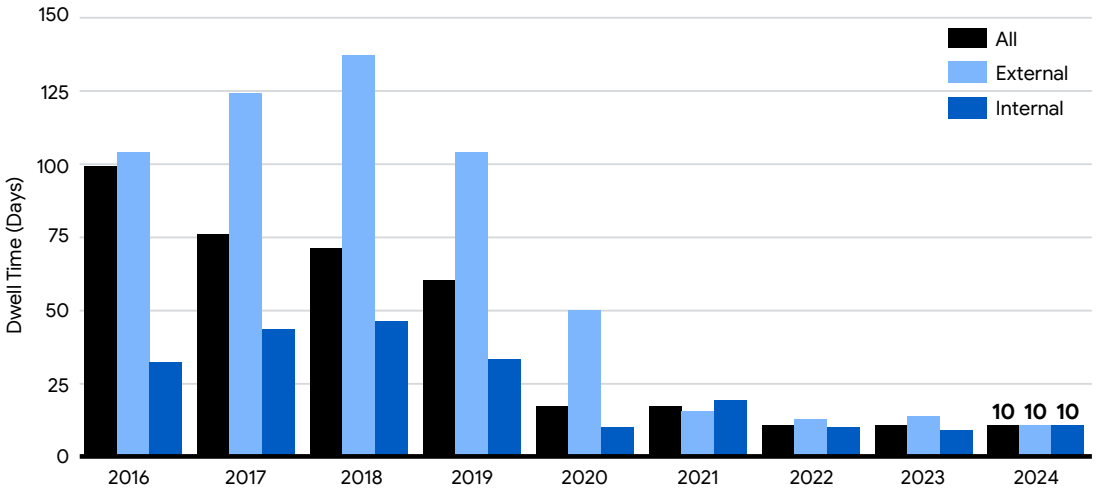

The dwell time distribution for the Americas in 2024 shows that, in aggregate, organizations continue to reduce the proportion of intrusions that remain undiscovered for long periods of time and increase the proportion of compromises that are discovered within a week of malicious activity. The percent of intrusions that lasted one week or less in the Americas in 2024 was 46.6%, compared to 45% in 2023.

Americas Dwell Time Distribution, 2021-2024

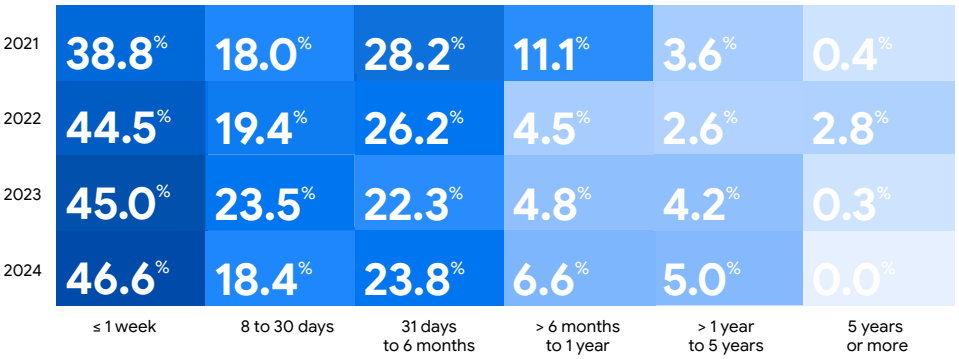

Threat Groups

The most frequently observed attacker in the Americas was UNC5267, which is the primary activity cluster Mandiant has designated to track North Korean IT workers. Mandiant responded to numerous intrusions involving North Korean malicious insiders who had applied to work at targeted organizations under false pretenses, misrepresenting their identities, locations, and legal status in order to generate revenue for the North Korean state.

The second most frequently encountered threat actor in Mandiant incident response investigations in the Americas in 2024 was the suspected Chinese cyber espionage actor UNC5221. The majority of observed activity was related to UNC5221 exploiting CVE-2023-46805 and CVE-2024-21887 in December 2023 and early 2024 to gain access to a number of organizations.

Mandiant investigators also identified UNC2565 at numerous investigations in the Americas in 2024. UNC2565 is a financially motivated threat cluster that uses the GOOTLOADER downloader to deliver a variety of secondary payloads, including BEACON, CLEANBOOST, LIGHTDUTY, SNOWCONE, and WORDFRAME. These intrusions have stemmed from victims accessing compromised websites. GOOTLOADER infections have been observed leading to data theft exfiltration and/or ransomware deployment.

EMEA

The metrics reported in this section are based on Mandiant Consulting investigations affecting organizations in Europe, the Middle East, and Africa (EMEA).

Targeted Attacks

Initial Infection Vector

The most frequently identified initial infection vectors in Mandiant incident response investigations in EMEA in 2024 were exploits (39%), followed by email phishing (15%) and brute-force attacks (10%). In EMEA, email phishing and brute-force attacks represented larger proportions of observed initial infection vectors than Mandiant encountered in global investigations.

EMEA

Exploit
39%

Email Phishing
15%

Brute Force
10%

Detection by Source

The intrusions that Mandiant investigated in EMEA in 2024 were first discovered internally 41% of the time, while in 59% of cases, an external organization first notified organizations of a compromise. These figures are similar to the global numbers (43% internal and 57% external).

EMEA Detection by Source, 2017-2024

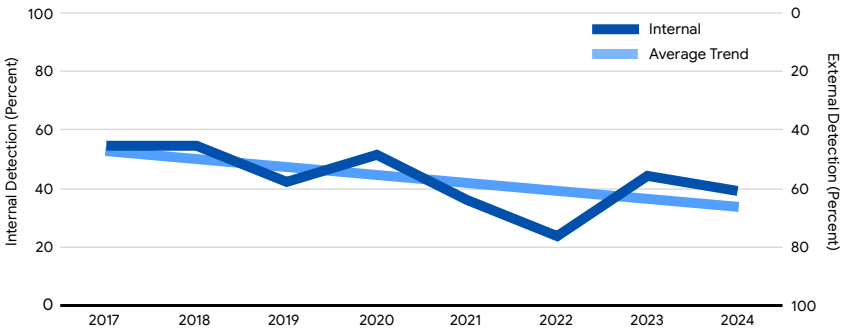

In contrast to the distribution observed globally, in Mandiant investigations in EMEA in 2024, adversary notifications comprised a relatively small share of notifications overall (3%) and ransomware-related events as well (12%). In all Mandiant investigations in 2024, adversary notifications represented 14% of overall incident discoveries, while adversaries notified organizations of a breach in 49% of ransomware-related events.

EMEA Detection by Source, 2024

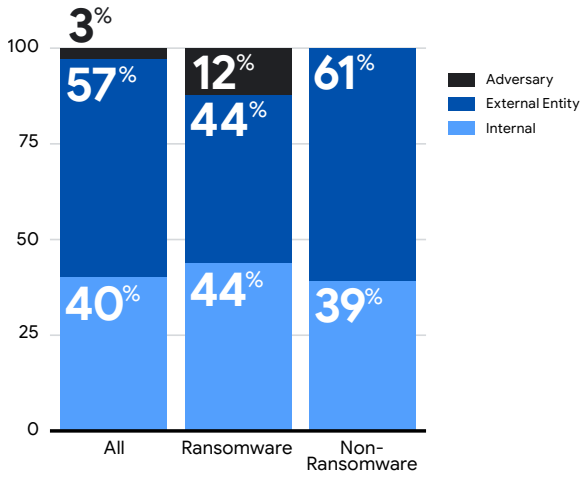

Median Dwell Time

The median dwell time for EMEA 2024 investigations was 27 days overall, 20 days for internally discovered events, and 32 days for externally notified events. While the 2024 median dwell times are higher than 2023 numbers for overall (22 days) and for externally notified events (12 days), over the long term, dwell times continue to decline. The median dwell time for ransomware-related events that Mandiant investigated in EMEA in 2024 was seven days, compared to 36 days for non-ransomware related intrusions.

EMEA Median Dwell Time, 2016-2024

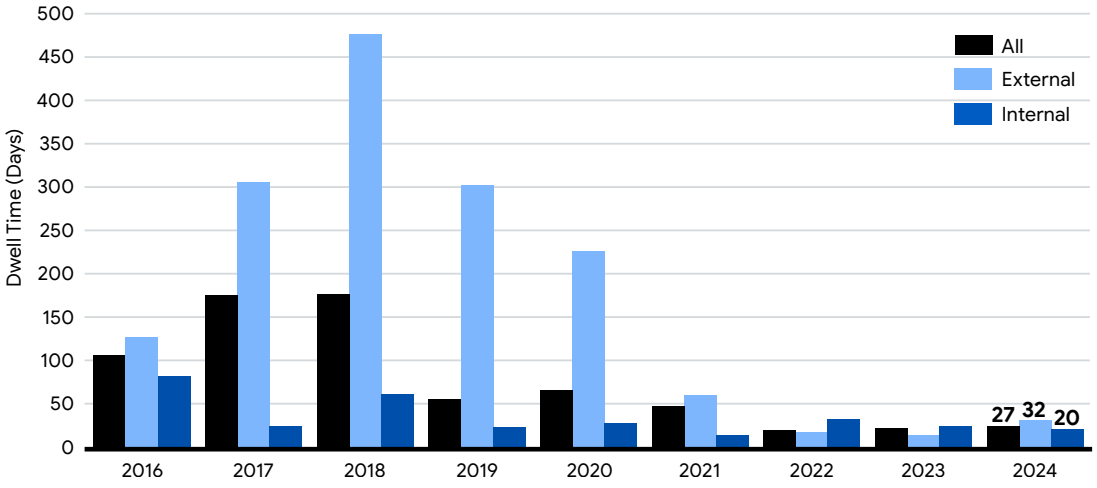

The dwell time distribution for Mandiant incident response investigations in 2024 in EMEA shows that the long-term trend is leading to fewer intrusions remaining undiscovered for long periods of time. The proportion of intrusions that were discovered within one week increased to 36.7% in 2024.

EMEA Dwell Time Distribution, 2021-2024

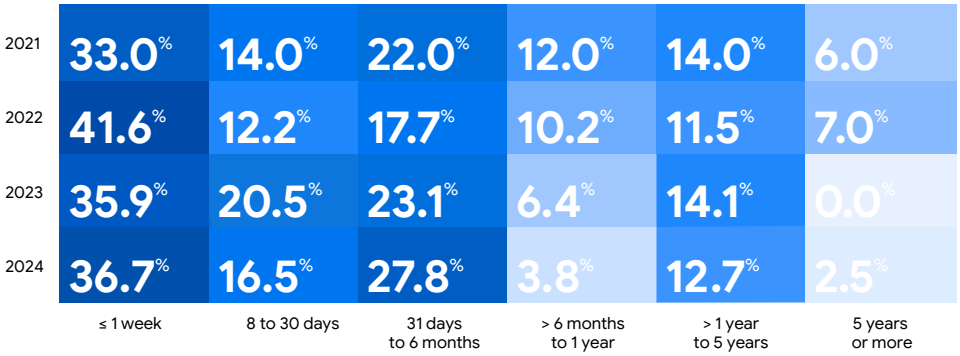

Threat Groups

Mandiant experts frequently encountered UNC4393 in 2024 investigations in EMEA. UNC4393 is a financially motivated threat cluster that has monetized access by deploying BASTA ransomware. In at least one case, Mandiant observed UNC4393 leveraging initial access established by a separate threat actor, UNC5155, using SILENTNIGHT malware. In other investigations, UNC4393 used brute-force attacks or stolen credentials to gain access to targeted environments.

In Europe, particularly in Ukraine, Mandiant continued to respond to APT44[22] intrusions in 2024. Mandiant believes that APT44 remains a core contributor to cyber operations related to the conflict and recently described how APT44 and other Russian cyber espionage threat clusters have demonstrated a focus on targeting mobile messaging applications for intelligence collection.

JAPAC

The metrics reported in this section are based on Mandiant Consulting investigations affecting organizations in Japan and Asia Pacific (JAPAC).

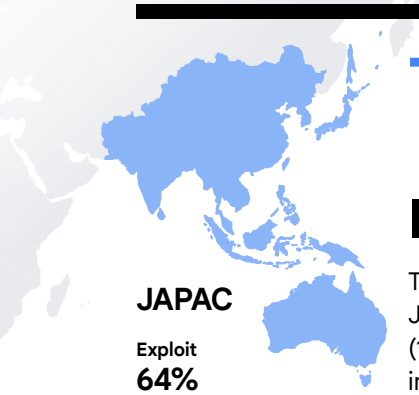

JAPAC

Exploit
64%

Stolen Credentials
14%

Web Compromise
7%

Targeted Attacks

Initial Infection Vector

The most frequently seen initial infection vectors in Mandiant investigations in 2024 in the JAPAC region, when they could be identified, were exploits (64%), followed by stolen credentials (14%) and web compromise (7%). Exploits and stolen credentials also topped the list for global investigations. Both in JAPAC and globally, use of stolen credentials eclipsed email phishing as an initial infection vector in 2024. The popularity of infostealer malware, as well as the widespread availability of credentials in data leaks and underground forums, may have contributed to increased incidences of this tactic. Organizations seeking to reduce exposure to the use of stolen credentials should ensure identity and access management policies that include multifactor authentication (MFA) are enforced across all user and account types.

Detection by Source

Organizations identified the first evidence of malicious activity internally in 31% of Mandiant investigations in the JAPAC region in 2024. External notifications accounted for 69% of detection sources. These figures are identical with detection sources for Mandiant investigations in the region in 2023.

JAPAC Detection by Source, 2017-2024

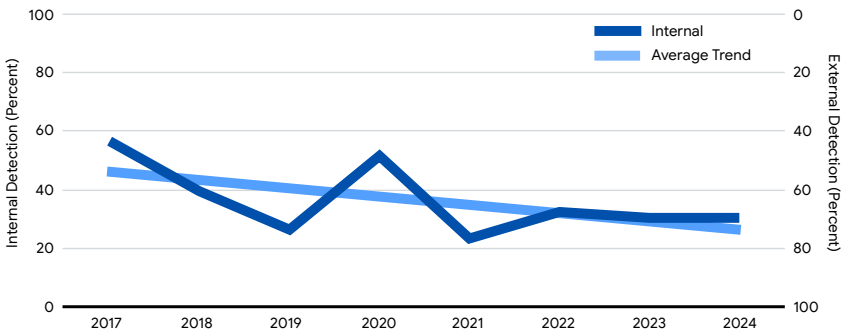

External notifications can also be divided into adversary notifications and external entity notifications from organizations such as law enforcement or cybersecurity companies. In 2024, Mandiant investigations in JAPAC, adversary notifications represented a smaller share of overall and ransomware-related events than in global numbers, with 12% adversary notifications in all investigations and 33% in ransomware-related intrusions, compared to 14% and 49% globally. External entity notifications for 2024 JAPAC investigations were proportionally higher than global numbers, at 57% overall compared to 43% globally.

JAPAC Detection by Source, 2024

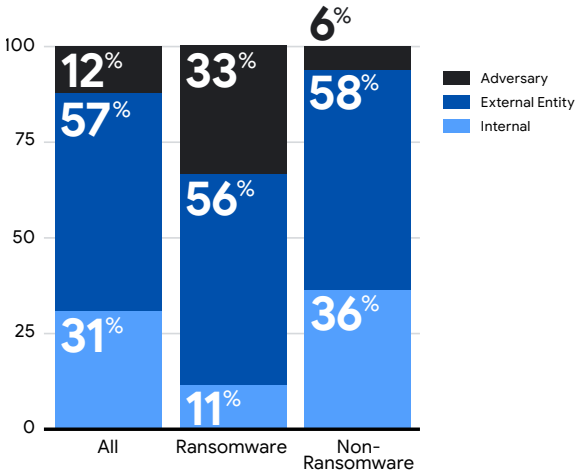

Median Dwell Time

The median dwell time for all intrusions in JAPAC in 2024 was six days overall, 10 days for externally notified events, and six days for internally discovered intrusions. For ransomware-related intrusions in JAPAC in 2024, the median dwell time was just four days. For non-ransomware-related compromises, the median dwell time increased to 12 days.

JAPAC Median Dwell Time, 2016-2024

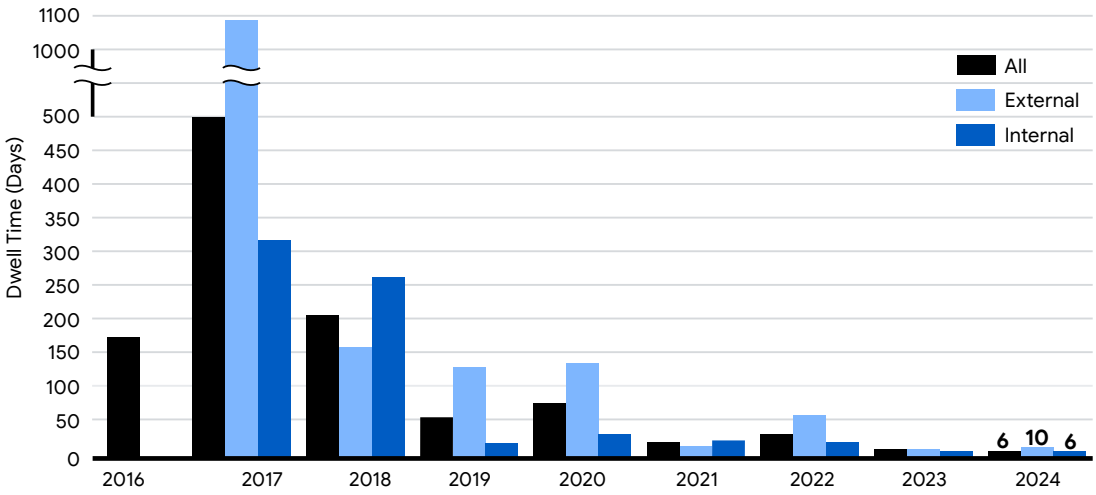

The dwell time distribution for JAPAC indicates incremental improvement each year in reducing the number of long-tailed compromises and increasing the proportion of malicious events that are discovered within the first week. In 2024, more than half of JAPAC investigations were identified within seven days of the first evidence of malicious behavior, an increase from 48.1% in 2023.

JAPAC Dwell Time Distribution, 2021-2024

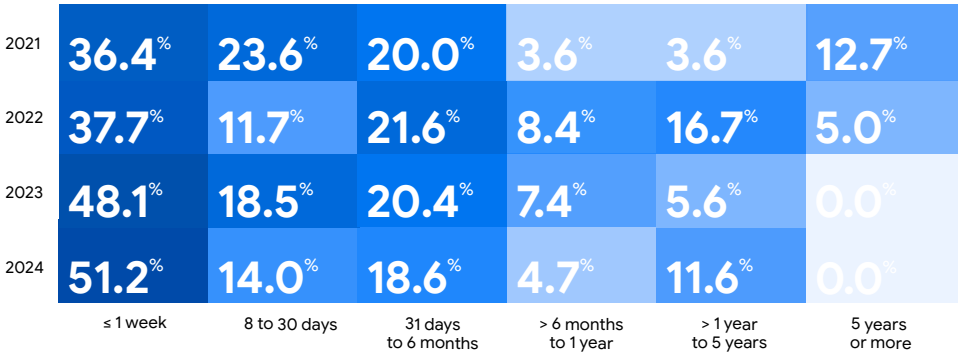

Threat Group

Mandiant incident response investigators identified UNC5221 activity during multiple engagements in JAPAC in 2024. UNC5221 is a suspected Chinese cyber espionage actor that exploited CVE-2023-46805 and CVE-2024-21887 in December 2023 and early 2024 to gain access to a number of organizations.

Articles

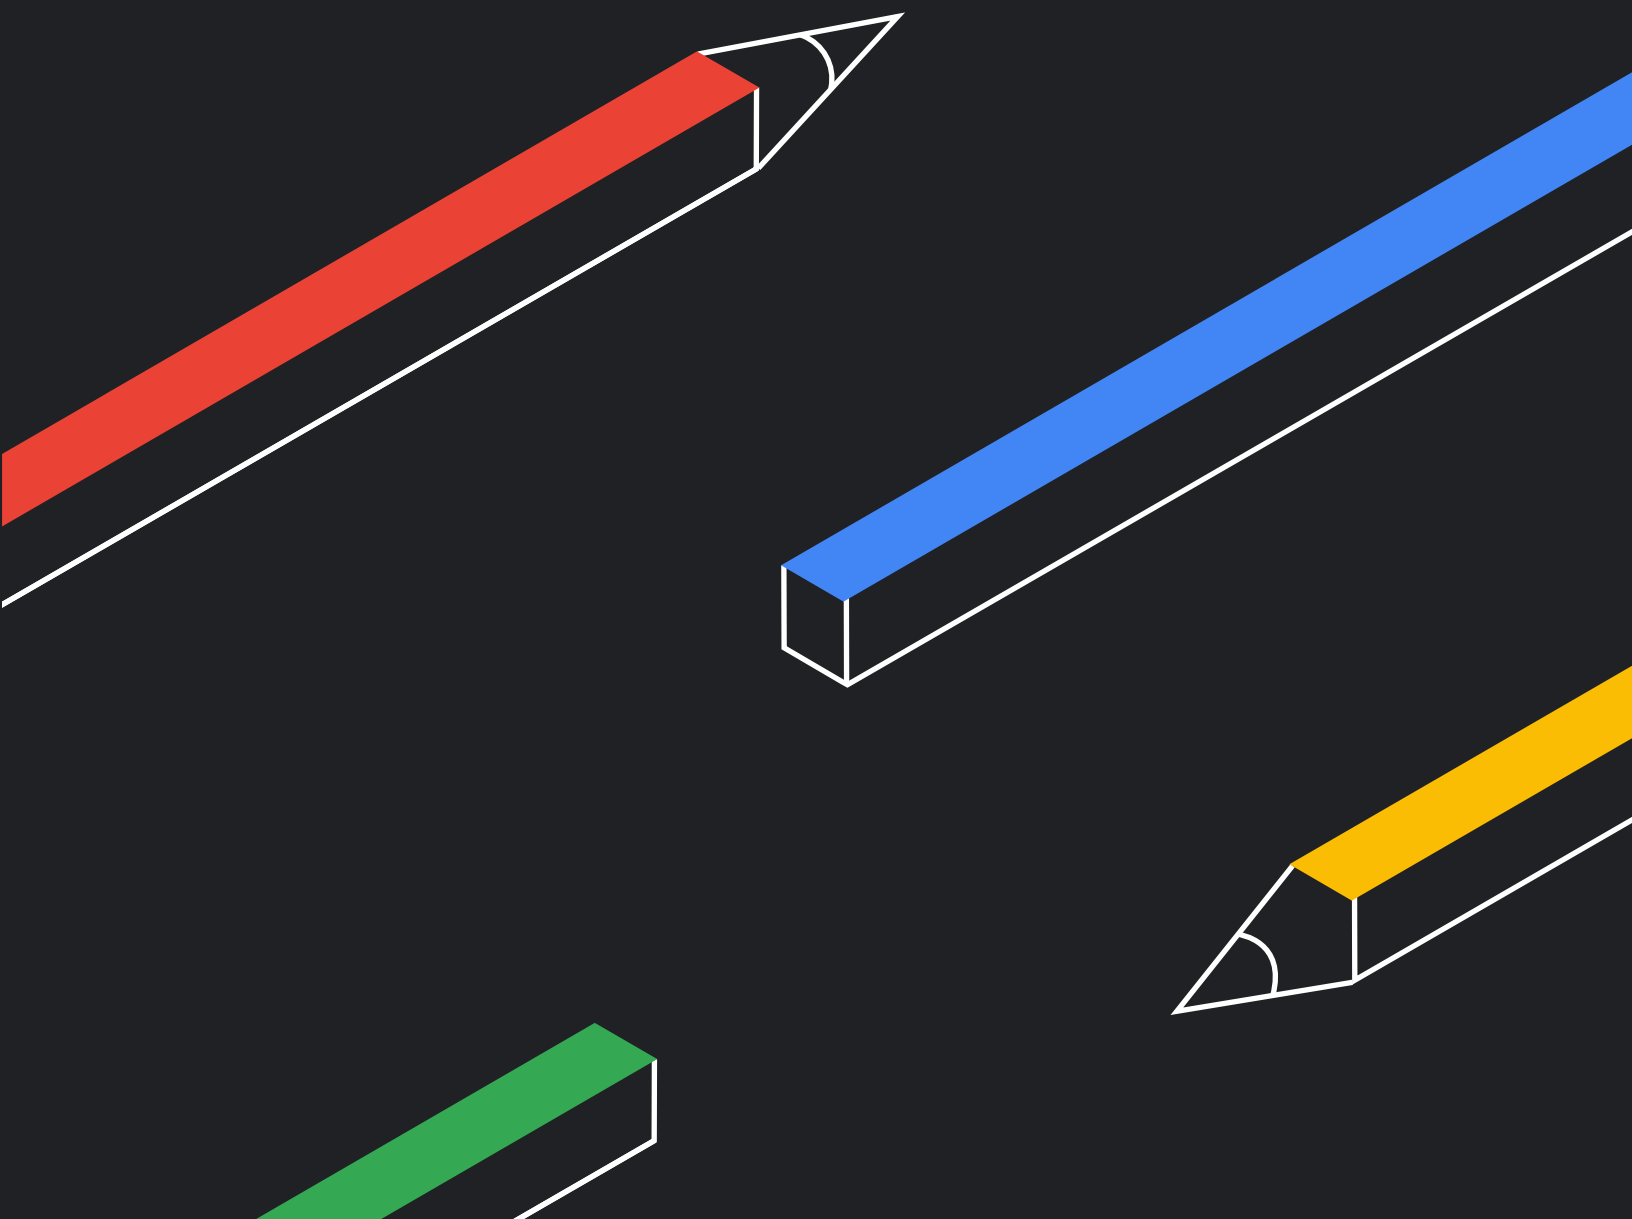

Infostealer Malware Continues to Create a Threat to Enterprise Systems

In the past several years, Mandiant has seen increased attention on a specific category of malware known as infostealers and their role in enabling often short-lived, yet deeply impactful intrusions using stolen credentials. Although infostealers and stolen credentials have always been a serious concern in cybersecurity, the recently renewed focus on infostealers by malicious actors and—consequently cybersecurity organizations—could signal drastic shifts in the ways cyber criminals abuse and/or monetize data obtained from infostealers.

Specifically, Mandiant has observed a resurgence in the use of stolen credentials as a means of initial access for compromises. While the use of stolen credentials by threat actors had dropped from 14% in 2022 to 10% in 2023, Mandiant identified stolen credentials in 16% of the intrusions observed in 2024. This resurgence is likely fueled, at least in part, by the large tranches of stolen credentials offered within cyber crime communities that have facilitated this rise in demand by offering stolen credentials in large tranches and on an individual basis.

Infostealers and broader credential theft are not new threats, but they are seeing a resurgence and have always posed significant risks to organizations that may not realize employee credentials have been compromised and exposed—sometimes years prior.

The Infostealer Problem

Infostealers are a broad classification of malware that have the capability of collecting and stealing a range of sensitive user information, such as credentials, browser data and cookies, email data, and cryptocurrency wallets. Notably, Mandiant does not classify malware used for mass data theft or collection of basic system survey information as infostealers. Examples of prominent infostealers include VIDAR, RACCOON, and REDLINESTEALER.

While many infostealers are built specifically for these purposes, they may also include basic backdoor and/or remote access trojan (RAT) capabilities, allowing them to be used to facilitate various attack lifecycle stages during intrusion operations. Further, infostealer capabilities can be added to traditional backdoors and RATs to extend the functionality of existing malware. For example, TRICKBOT, a malware family infamous for its use as a banking trojan and in intrusion operations, was also able to load a credential theft module for infostealing capabilities.

Information and credentials obtained via infostealers are commonly referred to as "logs" and are widely shared and sold across underground markets and criminal communities. Threat actors are able to search infostealer logs for information of interest specific to their targets. Infostealer logs can contain data that indicates the use of specific websites by users or even the specific software installed on the system. This information allows threat actors to more easily identify targets that align with the interests of their particular operations.

Mandiant has identified corporate credentials in infostealer logs, which highlight the risk to organizations. Successful compromise of an individual user could result in a threat actor gaining further access into an environment.

Example: UNC5537 Targets Snowflake Customer Instances for Data Theft and Extortion

Beginning in April 2024, a financially motivated threat actor, UNC5537, used stolen credentials to access the Snowflake customer instances of multiple organizations. These credentials were primarily obtained from infostealer malware campaigns that infected the work or personal computers of the employees and contractors that accessed Snowflake customer instances. This allowed the threat actor to gain access to the affected customer accounts and led to the theft of a significant volume of customer data from their respective Snowflake customer instances. Subsequently, the threat actor attempted to extort many of the victims directly and sought to sell the stolen customer data on cyber criminal forums.

Mandiant identified that the threat actor used Snowflake customer credentials[23] previously exposed via several infostealer malware variants, including VIDAR, RISEPRO, REDLINE, RACCOON STEALER, LUMMA, and METASTEALER. The earliest compromised credential leveraged by the threat actor was associated with an infostealer infection dating back to November 2020. In several Snowflake-related investigations, Mandiant observed that the initial compromise of infostealer malware occurred on contractor systems that were also used for personal activities, including gaming and downloading pirated software.

UNC5537's campaign against Snowflake customer instances was not the result of any particularly novel or sophisticated tool, technique, or procedure. This campaign's broad impact was the consequence of the growing infostealer marketplace, and it highlights the risk posed by the sheer volume of credentials circulating in these markets.

Unique Challenges of Infostealers for Enterprise Environments

Infostealers are often distributed broadly, typically targeting individuals, but they can also create unique challenges for organizations. Unlike other forms of credential theft, such as phishing and credential stuffing that can be used to target credentials for a specific system, infostealers can collect wide swaths of user data and credentials from a single host. Further, in cases where employees or contractors leverage personal devices for work purposes, the threat of infostealers can manifest outside of the scope of enterprise security and detection measures. For example, corporate credentials could be compromised when used on an infected personal device, or a compromised personal account could be leveraged in a password reuse attack against a corporate system. Browsers that support synchronization of passwords between instances can result in corporate passwords being synced to the personal systems of employees and may result in exposure. Policies to disallow and detect the use of browser syncing can help limit this exposure, especially when paired with user education, which trains employees to keep personal and corporate account use separate.

Contractors' devices, often used to access the systems of multiple organizations, present a significant risk. If compromised by infostealer malware, a single contractor's device can facilitate threat actor access across multiple organizations. In addition to being sold on underground markets, stolen credentials and information from infostealer infections are often shared openly in cyber criminal communities. This proliferation of infostealer logs and stolen credentials in these communities allows the information to remain available to threat actors indefinitely, where it can be used to impact organizations long after the infostealer infection occurred—in some cases years later.

A major advantage of obtaining accesses from infostealer logs is they can allow threat actors to search for specific types of accounts depending on their goals. The broad distribution of infostealers, coupled with the wide range of information they can collect from victims, provides a plethora of credentials and sensitive information for threat actors to work with. Accounts and services found in these logs, such as credentials for corporate virtual private networks (VPNs) and other enterprise services, can act as a foothold for further lateral movement within a network. Alternatively, actors may search infostealer logs for accesses tailored to other operations, including systems containing sensitive information for data theft extortion operations or cloud assets for illicit cryptocurrency mining activity.

Example: TRIPLESTRENGTH Leverages Stolen Credentials for Cloud Assets for Illicit Cryptocurrency Mining

Since 2023, teams across Google Cloud have worked to disrupt a financially motivated actor that the Google Threat Intelligence Group (GTIG) tracks as TRIPLESTRENGTH. This actor engaged in a variety of threat activity, including cryptocurrency mining operations on hijacked cloud resources. To take over cloud service accounts, TRIPLESTRENGTH leveraged stolen credentials and cookies to gain access to victim cloud environments. Once authenticated, the actor uses hijacked cloud projects to mine cryptocurrencies. Based on analysis of attacker-owned infrastructure, GTIG determined that the actor has relied on RACCOON infostealer logs as the source of at least a portion of the stolen credentials and cookies used in cloud hijacking activities and that the actor had access to credentials for Google Cloud, Amazon Web Services, and Linode. Additionally, in monitoring Telegram channels, Mandiant has observed personas connected to the group routinely advertise access to servers, including those from prominent hosting providers and cloud platforms, such as Google Cloud, Amazon Web Services, Microsoft Azure, Linode, OVHCloud, and Digital Ocean.

Recommendations

To mitigate the risk of infostealers, Mandiant recommends organizations leverage adversary-in-the-middle (AiTM)-resistant multifactor authentication (MFA) methods, such as hardware security keys or mobile authenticator apps. Organizations should consider implementing cookie expiration and password rotation policies to require regular password changes for accounts. This will limit the lifespan of any compromised credentials and cookies. Additionally, developing a robust access policy that restricts access from unknown or untrusted locations can limit threat actors' use of stolen credentials.

To further strengthen an organization's security posture against information-stealing malware, implementing endpoint detection and response (EDR) and intrusion detection systems (IDS) allows for fine-grained monitoring of environments. When configured and monitored effectively, these tools can provide comprehensive protection by detecting, preventing, and eradicating infections. As infostealers will commonly extract data from an end user's browser, organizations should apply controls to the browser to restrict third-party cookies, disable the use of autofill for passwords, and disable browser extensions that have not been approved for use.

To reduce the risk posed by external devices, such as personal devices, organizations should develop policies that strictly separate the use of personal and corporate systems. Organizations that rely on Bring Your Own Device (BYOD) should design policies or establish restrictions regarding appropriate use cases and ensure that additional measures, such as endpoint instrumentation for BYOD devices and MFA for passwords, are conditions that must be met for use. This will help to prevent malware threats from manifesting outside the scope of enterprise detections. Organizations should also review the security controls that third-party suppliers and contractors enforce on their devices to ensure malware threats from infostealers are not introduced via the supply chain.

Finally, infostealers are commonly distributed by disguising the malware as legitimate or cracked software. Organizations should establish software use policies and conduct training to prevent users from downloading software from untrusted sources. Organizations could also consider implementing an enterprise application store, where end users are empowered to download approved applications. IT security staff should validate these applications to ensure they are free from malware prior to being made available.

Detection Methods Based Around the Attack Lifecycle of Infostealers

Threat actors introduce infostealers using a variety of deceptive tactics. Phishing emails are a common method that involve using malicious attachments disguised as legitimate files or malicious links that lead to compromised websites or files hosting the malware. Compromised websites can also trigger drive-by downloads to automatically install the infostealer, sometimes using exploit kits to compromise browser or plugin vulnerabilities. Infostealers may also be bundled with infected software downloads from untrusted sources or included in trojanized versions of legitimate software. Finally, attackers use social engineering to manipulate users into downloading or installing the malware.

To prevent infostealer infections upon initial delivery, organizations should use existing security infrastructure to analyze network traffic and email. Email gateway monitoring can flag suspicious emails that bypass initial filters, enabling further review and potential intervention. Additionally, monitoring outbound network traffic via proxies and intrusion detection systems, as well as reviewing DNS requests, can help detect malicious downloads. Most enterprise firewalls, DNS servers, and proxies offer built-in monitoring capabilities. Ensuring these detections are sent to a security information and event management (SIEM) platform and reviewed by a security team is a crucial step in limiting the spread of infostealers. If these events are not investigated, malware may be detected but not properly remediated if the infection bypasses EDR and antivirus.

Infostealers often evade antivirus and EDR tools by manipulating system resources and behaviors. For instance, dynamic-link library (DLL) side-loading takes advantage of the Windows loading process to substitute malicious DLLs for legitimate ones, thereby hijacking application functionality. They may also disable or modify security tools, either by altering configurations or outright disabling them. To further conceal their presence, infostealers sometimes use hidden files and directories, complicating malware analysis and identification.

Conclusion

While infostealers and broader credential theft are not novel techniques, we anticipate that actors of varying motivations and levels of sophistication will continue to demonstrate a significant interest in leveraging stolen credentials as an initial intrusion vector. Infostealers can be an effective method for obtaining stolen credentials as they are capable of collecting wide swaths of user data, are readily available in underground communities, and allow actors to easily search logs for special accesses of interest. Given the wide availability and long-standing presence of infostealers in underground communities and illicit operations, organizations must be aware of the direct and indirect risks posed by infostealers.

Democratic People's Republic of Korea Insider Threats

Due to international sanctions placed on the country in 2003, the Democratic People's Republic of Korea (DPRK) has sought to identify means through which they can continue to fund national interests. As sanctions intensified in 2016 in response to the DPRK's testing of nuclear weaponry and as a means to further impact the ruling class in North Korea, the country found itself cut off from financial systems in the West, further limiting its ability to generate revenue. In response, the DPRK has pursued a variety of means to evade sanctions, including illegal weapons sales, front companies operating in international regions, and outright theft. As technology has progressed, the revenue-generating schemes pursued by the DPRK have evolved. Ranging from the theft of more than $100 million USD through fraudulent SWIFT transactions in 2016 to compromises targeting cryptocurrency in 2024,[24] technical proficiency leveraged for theft has been a primary focus for the DPRK.

Since 2022, Mandiant has tracked a threat cluster it refers to as UNC5267, which represents the DPRK's efforts to place thousands of its citizens in countries outside of North Korea to pose as remote IT contractors for Western companies. These citizens, commonly referred to as "DPRK IT workers," are directed to seek employment in high-tech companies headquartered among Western countries and funnel salaries back to the DPRK to fund national interests, including the continued investment in weapons of mass destruction. DPRK IT workers most commonly work through job placement services and recruiters but have been observed pursuing direct employment as well. Their operations are supported through a broad network of false or stolen identities and third-party accomplices. Outside the fraudulent activity necessary to place a DPRK IT worker in a Western organization, Mandiant identified evidence of direct malicious activity in fewer than five investigations in 2024. However, the access to corporate infrastructure necessary for the high-tech jobs that DPRK IT workers pursue places organizations at heightened risks of extortion, espionage, data theft, and disruption, which may escalate as the campaign continues.

Pre-Hiring Tradecraft

The fraud and identity theft guardrails surrounding employment in Western countries require both applicants and employers to adhere to a strict set of processes designed to limit the hiring of individuals using fraudulent identities. As such, long-term employment of a North Korean citizen with the ultimate goal of funneling money back to the DPRK without exposure requires the creation of a complex network of false personas and supporting documents. DPRK IT workers have been observed using stolen identities and identities that appear to be wholly fabricated to support their operations. Each persona and supporting document—or element of a falsified online presence—comes with its own care and feeding requirements to maintain the illusion of a potential dedicated employee. Similarly, the language requirements needed to navigate the interview process successfully can add an additional strain on the upkeep of the false persona in use. While the DPRK has invested heavily in education for the English language, science, and math, maintaining what is effectively a cover identity in a foreign language for a single identity is a taxing endeavor. For DPRK IT workers, however, it appears they maintain a substantial array of false identities.

The competitive nature of the IT industry makes individual efforts at placing a DPRK IT worker in a high-paying position far from guaranteed. In 2024, Mandiant identified a suspected DPRK IT worker using at least 12 personas while seeking employment in the US and Europe. DPRK IT workers have been observed providing references to recruiters for other false personas controlled by the DPRK. In at least one instance, two false identities were considered for a job in a US company, with one DPRK IT worker winning out over the other. In at least three investigations, Mandiant identified multiple suspected DPRK IT workers hired by the customer. In one such example, four suspected DPRK IT workers had been employed within a 12-month period at a single organization. Successfully navigating an organization's hiring process may give DPRK IT workers adequate experience such that they can continue to target that organization using additional false personas.

Mandiant has identified suspected DPRK IT worker profiles hosted on job-posting platforms such as LinkedIn and Indeed that contain false testimonies, fabricated employment and educational histories, and which claim a wide range of technical proficiency. Online profiles maintained by suspected DPRK IT workers are often carefully crafted, with some even going so far as to interact with officers of the universities from which they claim to have graduated. A pattern commonly found on resumes of suspected DPRK IT workers is one in which the persona claims to reside at a local US-based address but to have studied abroad at international universities. This pattern is not wholly consistent across all suspected DPRK IT worker profiles but may serve to hinder the efforts of potential employers seeking to confirm the educational background of a false persona. Similarly, when an applicant undergoes a background check during the interview process, DPRK IT workers have been observed providing education histories that do not match the program of study or the years of attendance listed on their resume.

Much like any organization facing administrative burdens, DPRK IT workers have found they can alleviate some of the overhead through simple reuse. Resumes associated with suspected DPRK IT workers can be seen borrowing heavily from publicly available resumes, and even reusing those among the corpus of resumes for DPRK IT workers. Mandiant's analysis of a Netlify page associated with a suspected DPRK IT worker[25] uncovered two distinct resumes that presented separate identities with unique personal information, such as phone numbers and email addresses. The resumes listed differing educational and professional backgrounds, but both included identical uncommon phrases, which could be used to tie the resumes to a potential singular author. The supporting sites, which are used to bolster the false persona used by a DPRK IT worker, also appear subject to a degree of "templatization." Sites suspected to be part of DPRK IT worker operations often reuse common themes, layout, and content or leave key sections unaltered from their defaults.

Among the content found on suspected DPRK IT worker sites, resumes, and postings, Mandiant has identified additional patterns of use for various key artifacts. Email addresses and domains commonly include a series of themes, including specific words or numbers. The words "panda," "dev," "star," "silver," and "sun" are often reused across a series of indicators associated with DPRK IT workers. In an affidavit filed by the US Federal Bureau of Investigation (FBI) in 2023,[26] the FBI attested that a freelance worker who provided account sharing to a suspected DPRK IT worker was supplied with a security challenge password that, when translated from Chinese, referred to "silver star." The reuse of key artifacts in background material for false personas reduces the level of effort needed to maintain a variety of ready-to-use identities containing a mixture of fabricated data often overlaid atop identities stolen from US citizens.

Post-Hiring Tradecraft

DPRK IT workers have been reported as residing primarily in Russia and China,[27] with smaller groups suspected to reside in Africa and Southeast Asia. Geographical location has long been a reliable means for detecting fraudulent activity across many security realms. While network traffic originating from North Korea would raise immediate alarm bells in an organization's security operations center, the countries most accessible to DPRK agents, such as Russia and China, share the same threat characteristics for many Western organizations. Even for organizations that might not alert over simple geographical associations, the disparity between the location of the falsified persona and the region from which their connections originate expose additional risk for detection. To reduce the risk of exposure, once engaged with an unsuspecting employer, DPRK IT workers rely on a variety of techniques to maintain operational functionality while obfuscating their identity and location.

Since Mandiant began tracking DPRK IT worker activity in 2022, Mandiant has observed suspected DPRK IT workers connect through virtual private network (VPN) sessions associated with the Astrill VPN in 72% of investigations. Threat actors across all levels of complexity and motivations have recognized that a VPN can raise the level of effort required for network defenders to identify potentially malicious network sessions effectively. This can be as simple and useful as threat actors using services that terminate in a Western country, while the threat actors themselves operate from a country that would appear more suspicious. To combat this kind of threat among a growing remote workforce, many companies rely on impossible travel analysis to identify sessions that indicate a user has connected from a region they could not have traveled to in the time between successive connections. Mandiant has observed advanced threat actors going as far as to ensure their connections originate from the same region as a legitimate connection would originate for a specific compromised user. During such incidents, investigators must work to differentiate the expected activity from malicious activity among multiple sessions occurring in close geographical proximity. For suspected

DPRK IT workers, however, VPN analysis is further complicated due to the kinds of work the operative engages in on behalf of their employer and an overall lack of malicious activity. In cases involving suspected DPRK IT workers, the actions taken rarely, if ever, step into the category of malicious activity commonly associated with threat actors. Instead, their activity blends into legitimate network traffic almost entirely.

Since the wide adoption of remote work, provisioning and shipping a corporate laptop to newly hired remote workers has become a common onboarding process for many organizations. This provides organizations more control over the individual systems that connect back to the corporate environment. Similarly, security teams have an opportunity to apply policies and instrumentation to the endpoint, which grants a greater degree of visibility and the ability to limit the specific applications allowed on the endpoint. This model for onboarding new hires introduces an additional avenue of risk for DPRK IT workers as shipping a laptop to their physical locations is likely to raise an immediate alarm within the organization. This has led suspected DPRK IT workers to rely on in-country "facilitators" who perform services for a fee. Facilitators supporting DPRK IT workers have been identified in the US and in Europe. The services provided by facilitators range from simple singular interactions to contracts with an expectation of ongoing support.

In some cases, facilitators may assist with cashing paychecks or receiving physical mail, including corporate hardware on behalf of their customers. In one case, a facilitator was used to pass an in-person drug test for a DPRK IT worker hired by a US company. Mandiant investigated a suspected DPRK IT worker compromise in 2023 during which the operative's corporate laptop was shipped to an apartment block in a major US city. When law enforcement investigated the location, they found an empty apartment and the box in which the laptop was shipped. An analysis of connection logs indicated that the suspected DPRK IT worker connected to corporate resources through an Astrill VPN session that masked the origin of the connection. A subset of the network sessions recorded showed the VPN connection appeared to fail and, before being reestablished, IP addresses associated with China were observed connecting to the same corporate resource from the same system. A facilitator was used in this case to receive and potentially reship the laptop from its expected location to the true location of the remote DPRK IT worker. In this instance, the customer had not suspected their employee of operating under a false persona until notified by law enforcement.

On the other end of the spectrum of support provided by facilitators, some operate full "laptop farms," which host the corporate laptops of their customers for remote access. This provides a stable location from which network connections will be sourced, which matches the country in which the company is headquartered. Facilitators ensure laptops remain active and available for their customers and install remote access software that their customers use to access the corporate laptop. In two separate investigations performed by Mandiant incident responders, suspected DPRK IT workers provided the same shipping address for their corporate laptop during onboarding. A US grand jury indictment filed in 2024[28] against a suspected facilitator estimated that the accused knowingly assisted in fraud schemes, including running a laptop farm, which ultimately impacted more than 300 US companies using over 60 stolen identities, and resulted in at least $6.8 million USD in revenue for the DPRK.

While preconfiguring a system and onboarding remote users with a corporate-owned laptop is a common process, ongoing monitoring and restrictions on unnecessary applications is a less consistent operation. Remote management tooling is common in both legitimate and malicious use. DPRK IT workers are not unique in their understanding that legitimate remote access management tools have as much or more value in maintaining long-term access to an environment. While malware such as backdoors might provide more features, for threat groups pursuing more clandestine operations, the use of remote management tools reduces their detectable footprint. Much like the use of location-specific VPN sessions, DPRK IT workers enjoy a substantially reduced detection footprint as their day-to-day workflows are often indistinguishable from those of legitimate employees.

Detection and Mitigation

Detecting potential DPRK IT workers requires a strict employee data verification pipeline and a comprehensive baseline of endpoint and network monitoring. The best means through which organizations are able to protect themselves from this kind of risk are preventative measures. Additional scrutiny in the hiring process and improvements in the overall instrumentation and monitoring post-hiring are also valuable tools for organizations employing remote workforces.

During the interview process and when onboarding new remote workers, identifying disparities between the purported facts and the observed facts grants organizations an opportunity to protect themselves. Strict

background checks that include the collection of biometric information from the new hire, which is subsequently used for specialized background checking services, may help detect forged identities. Even identifying the service associated with the applicant's phone numbers could be a valuable check in the interview and onboarding process. Mandiant has observed suspected DPRK IT workers supplying phone numbers associated with Voice over Internet Protocol (VoIP) services instead of consumer phone lines. Similarly, logging and reviewing key artifacts, such as the email address and phone number used by applicants, can help develop a dataset against which hiring organizations can compare current and previous applicants to identify someone trying to reuse information under a different identity.

DPRK IT workers have often demonstrated an unwillingness to appear on camera during interviews and once hired. Differences in the personas they adopted in order to be interviewed, especially when using stolen identities, may become apparent to hiring managers and coworkers when they are forced to appear on screen. Rescheduling or outright cancelling interviews with candidates for remote work who refuse to appear on screen raises the burden for potential DPRK IT workers during the interview process. Forcing operatives to match their false personas to a specific physical presentation or rely on unproven technology such as video face-swapping services to bypass immediate detection also increases the chances they are detected by external security organizations and law enforcement.

Many of the suspected DPRK IT worker cases Mandiant investigated in 2024 stemmed from notifications provided to impacted organizations by law enforcement organizations, such as the FBI. Once hired, the detection opportunities available to organizations rely less on comprehensive security practices and more on identifying inconsistencies and exploiting mistakes made by suspected DPRK IT workers. Disparities between the geographical region in which a suspected DPRK IT worker purports to live and the addresses provided for shipping documents and corporate resources provide another opportunity for detection. DPRK IT workers have been observed using stolen identities and falsified identification credentials that retain the address of the original identity. In such cases, they often request corporate resources to be shipped to the address of an in-country facilitator. Requiring in-person pickup of corporate laptops with full verification based on a valid ID limits the ability for DPRK operatives to ship corporate hardware to laptop farms or to a follow-on destination. In the event that a remote hire requests corporate resources be sent to an address not listed on their employment documents, delaying shipment and reviewing the associated background checks may help reduce the hiring organization's exposure to risk.

From a technical standpoint, ensuring corporate resources are delivered with monitoring tools such as endpoint detection and response tooling pre-installed helps organizations build baseline application use metrics. Monitoring solutions should be configured to identify and alert on the use of remote access software and connections originating from VPN services. Endpoint detection and monitoring solutions should be configured to log details of any human interface devices (HID) plugged into the laptop, and this data should be reviewed. Mandiant has observed DPRK IT workers and facilitators use network-based KVM switches to control corporate laptops housed within laptop farms. Reviewing HID connect and disconnect logs is a crucial opportunity to identify potential DPRK IT workers.

Appropriately siloing data and conforming to a security framework that enforces the principle of least privilege should be a standard part of an organization's security posture. Ransomware operators, insider threats, and espionage groups all rely on access to data that exceeds what is necessary for most corporate roles. While evidence of direct malicious activity has been limited, in at least two cases Mandiant investigated in 2024, the suspected DPRK IT worker resorted to extorting their employer after they were exposed. In both instances, the exposed employees demanded money in exchange for promises to not publish confidential corporate data. Ensuring users only have access to the data and resources needed to perform their duties helps limit impact from a variety of threats, including those posed by DPRK IT workers.

Conclusion

The organizations DPRK IT workers target appear to align more with opportunistic targeting than with a given targeting objective. Additionally, the limited instances of direct malicious cyber activity point more toward targeting of high-paying job roles. One of North Korea's primary strategies for avoiding the negative effects of international sanctions is by finding ways to generate revenue. Furthermore, the continued pursuit of weapons of mass destruction is a primary goal of the Kim regime, with ever-growing budgetary demands. A large portion of suspected DPRK IT workers are reportedly subordinate

to organizations under the 313 General Bureau of the Munitions Industry Department,[29] which is responsible for the nuclear program in North Korea.

Organizations outside North Korea are natural targets for DPRK-nexus threat actors, either through the data they produce and store or by the simple fact that they generate revenue that can be funneled into the DPRK illicitly. The DPRK IT workers are the latest in a long series of tactics undertaken by a regime that is focused on evading international punitive measures. If not curtailed, DPRK IT workers operating within Western organizations pose a significant risk to businesses and national security beyond simple fraudulent employment.

The 2024 Iranian Threat Landscape

As tensions in the Middle East escalated throughout 2024, Mandiant observed the scale of Iran-nexus threat actor operations increase across the region. Iran-nexus threat actors continued to sustain cyber operations against targets of strategic and operational relevance, while increasingly focusing on Israeli targets.

Mandiant observed Iran-nexus threat actors combine several approaches to heighten the likelihood of successful intrusions. Most notably, they significantly expanded their arsenal of custom malware for use in the full spectrum of cyber operations. At the same time, they also maximized their use of publicly available resources such as cloud infrastructure and legitimate tools to evade detection. Mandiant observed threat actors employ increasingly effective social engineering schemes that quickly integrated worldwide events, computer security incidents, and employment themes. This resulted in effective campaigns through which Iran-nexus threat actors pursued cyber operations in alignment with national and strategic objectives.

Expanding Arsenal of Custom Malware

When conducting cyber operations, threat actors can choose to create their own malware or use readily available public tools. Proprietary malware allows the threat actor to tailor the malware to operational requirements. However, the flexibility provided by custom malware comes at the cost of resource-intensive development and maintenance. In the event that custom malware is discovered, replacing the capability can be costly for threat actors. In comparison, publicly available tools are more easily replaceable but may not fit all the threat actor's needs. Mandiant observed Iran-nexus threat actors align with the first approach throughout 2024 as they significantly increased their arsenal of custom malware.

Mandiant tracked a 35% surge in malware attributed to Iran-nexus threat actors compared to 2023, with more than 45 new malware families discovered in 2024. This increase may be due, in part, to the escalation of geopolitical tensions as a result of Iran's proxy war with Israel and its steady investment in offensive capabilities as part of a broader strategy to enhance cyber operations.

Destructive and Disruptive Malware

In 2024, Israel-based targets were a focal point for destructive and disruptive operations from Iran-nexus and pro-Iran threat actors. During these campaigns, Iran-nexus threat actors relied heavily on wipers, a type of malware designed to erase or corrupt the data of the computer it infects. While organizations associated with Israel were heavily targeted, entities in other regions, such as Albania, were also targeted with wiper malware early in the year by the same groups. These campaigns are often coordinated with exposure efforts by online personas with the ultimate goal of manipulating the public narrative surrounding regional issues. Ongoing hack-and-leak operations from various online personas affiliated with Iran-nexus threat actors aided in this endeavor.

The online personas that support disruptive operations often operate under the guise of cyber activism, also known as hacktivism, in an attempt to hide their affiliation with state-level entities. The online personas "Karma" and "Homeland Justice" have claimed credit for operations targeting organizations in Israel and Albania. In 2022, Homeland Justice claimed credit for an attack targeting Albania with the ROADSWEEP[30] malware. In 2023 and 2024, Karma claimed credit for wiper attacks[31] on Israeli organizations. Public reporting has asserted that both of the online personas were provided access to their targets through prior compromises from UNC1860,[32] which is publicly referred to as "Sacred Manticore." UNC1860 is likely affiliated with the primary intelligence agency of Iran, the Ministry of Intelligence and Security (MOIS).

Similarly, the online persona "Handala Hack" claimed responsibility for numerous cyberattacks that targeted Israeli government and financial organizations with the proprietary COOLWIPE wiper in December 2023. In July 2024, Handala Hack claimed responsibility for a phishing campaign that deployed COOLWIPE to Israeli targets. A more recent campaign delivered malware masquerading as a security patch for a faulty security vendor update. However, to date, evidence supporting the claims made by Handala Hack has not been provided.

The online group "Cyber Toufan" has also been linked to wiper activity and claimed responsibility for hack-and-leak operations targeting Israeli companies, government entities, and individuals. On the one-year anniversary of the Oct. 7 attack against Israel, Cyber Toufan promoted a video on their Telegram channel that corresponded with an operation that targeted Israel-based users with the proprietary POKYBLIGHT wiper. The same group was linked to Android and Windows wiper campaigns targeting Israel-based users earlier in 2024. In each instance, the phishing emails masqueraded as alerts to security fixes and safety guidelines from an Israeli government institute.

Proprietary Malware

Mandiant identified more than 20 proprietary malware families—including droppers, downloaders, and backdoors—used in campaigns in the Middle East. Six previously unknown custom malware families were deployed in 2024 as part of suspected APT34 operations targeting Iraqi government entities. APT34 is an Iran-nexus cyber espionage group that has been operational since at least 2014 and has been largely focused on phishing efforts to benefit Iranian nation-state interests. Two of the six newly identified backdoors, DODGYLAFFA and SPAREPRIZE, overlap with a public report[33] of suspected Iranian operations targeting Iraqi government networks.

UNC3313,[34] an Iran-nexus threat group that carries out surveillance and strategic information-gathering operations, was observed distributing a series of custom dropper and backdoor malware during spear-phishing campaigns in 2024. The threat actor hosted malware on popular file-sharing services and embedded links within training- and webinar-themed phishing lures. In one such campaign, UNC3313 distributed the JELLYBEAN dropper and CANDYBOX backdoor to organizations and individuals targeted by their phishing operations. UNC3313 is suspected to be affiliated with MuddyWater, a group the US Government reported[35] as being subordinate to the MOIS.

Prevalence of Graphical User Interfaces in Malware

In 2024, Mandiant observed an increased focus on deception techniques used to improve the chances of success when targeting individuals. Iran-nexus threat actors incorporated graphical user interfaces (GUIs) to disguise malware execution and installation as legitimate applications or software. The addition of a GUI that presents the user with a typical installer and is configured to mimic the form and function of the lure used can reduce suspicions from targeted individuals.

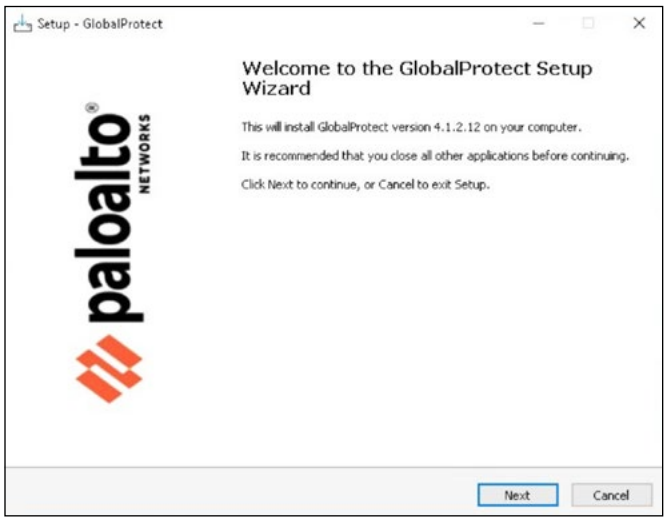

Figure 3.1: Wizard installation window displayed to the victim

In July 2024, a suspected Iran-nexus threat actor distributed the CACTUSPAL backdoor,[36] which masqueraded as an installer for the Palo Alto Networks GlobalProtect remote access client. Upon execution, an installation wizard that mimicked a legitimate Palo Alto Networks installer was displayed to the user while CACTUSPAL's .NET payload was written to disk. Once the targeted user closed the dialog window, the GUI thread aborted, and the main CACTUSPAL execution continued. The CACTUSPAL backdoor is designed to verify that only one instance of the process is running when executed before it initializes the staging directory and running configuration prior to the start of command-and-control (C2 or C&C) activity.

UNC2428, an Iran-nexus threat actor that conducts cyber espionage-related operations, is suspected to have distributed the MURKYTOUR backdoor through a complex chain of deception techniques in October 2024. UNC2428's social engineering campaign targeted individuals while posing as a recruitment opportunity from Israeli defense contractor, Rafael. Individuals who interacted with the campaign were redirected to a site purporting to be part of Rafael's web presence, where users could download a tool to assist with applying for a job. The installer, named RafaelConnect.exe, was the LONEFLEET installer malware, which presented the user with a GUI front-end through which they could provide personal information and an opportunity to submit a resume. After the form was submitted, the MURKYTOUR backdoor was launched as a background process. UNC2428's activity overlaps with the Israel National Cyber Directorate's attribution to a group called "Black Shadow."[37]

Figure 3.2: Suspected Black Shadow attack flow

Leveraging Cloud and Public Resources to Evade Detection

While Iran-nexus threat actors have invested in developing custom malware in recent years, they have also taken steps to reduce the detectable footprint of their intrusions. Mandiant observed Iran-nexus threat actors adopt greater use of legitimate remote monitoring and management (RMM) tools and tailor their operational infrastructure to mimic those used by their targets.

RMMs are legitimate tools that allow IT personnel to access a computer remotely in order to manage the system on which the tool is installed. UNC3313 relied heavily on RMMs during the initial access phase of many of their intrusions in 2024. Mandiant identified at least nine different RMM agents disseminated by UNC3313 in phishing campaigns over the year. During these campaigns, the threat actor would host the installer for a given RMM on major file-sharing services, with links to the installers included in various phishing lures. Upon installation, the RMM was configured to provide access to the system from attacker-controlled infrastructure. Since the RMMs used by UNC3313 had legitimate use cases, the likelihood of detection by network or endpoint agents was reduced when compared to a custom backdoor. Where a threat actor's use of custom malware can be exposed and quickly integrated into blocklists or endpoint detection and response (EDR) tooling, RMMs are rarely included in automated detect-and-block mechanisms due to the nature of the tools themselves. This can lead to a much delayed response between identification and actioning within an organization.

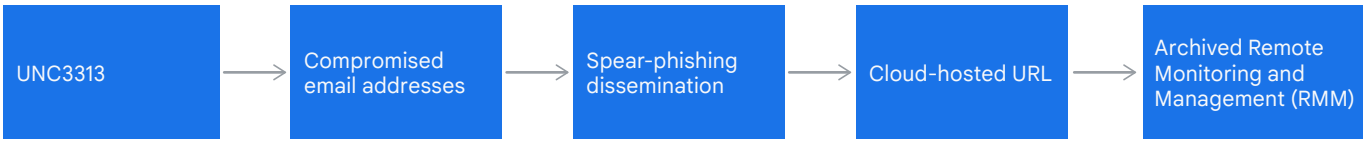

Figure 3.3: UNC3313 attack flow

A number of Iran-nexus threat actors have also been observed taking additional steps to ensure the infrastructure used during their attacks blended in with commonly used infrastructure. As cloud adoption continues to grow year-over-year, threat actors have taken advantage of the centralization of resources among the major cloud vendors. In addition to techniques such as typosquatting and domain reuse, threat actors have found that hosting C2 nodes or payloads on cloud infrastructure and using cloud-native domains reduces the scrutiny that may be applied to their operations.

UNC1549,[38] a suspected Iran-nexus threat actor that has targeted the aerospace, aviation, and defense industries in Middle East countries, regularly used cloud infrastructure during intrusions in 2024. The threat actor is suspected to have built C2 infrastructure and hosted payloads in the cloud while also tailoring the domain names they used to match common domains. In some cases, UNC1549 customized the domains used in their campaign on a per-target basis, and in others, went so far as to ensure servers were geolocated near their targets.

APT42,[39] a prolific Iran-nexus threat actor known for its meticulous social engineering efforts and rapport building, maintained a series of credential harvesting campaigns in 2024. Active since at least 2015, APT42 commonly maintains contact with targeted individuals as they attempt to build trust; they also often build well-tailored decoy sites during campaigns. Mandiant recently observed the threat actor deploying fake login sites mimicking Google, Microsoft, and Yahoo as part of their

credential harvesting campaigns. APT42 used cloud-based platforms and services, such as Google Sites and Dropbox, in operations that directed targets to fake Google Meet landing pages or login pages. The threat actors also targeted Israel and the US[40] in 2024, including individuals affiliated with presidential campaigns, military personnel, diplomats, academics, and non-governmental organizations (NGOs). APT42 deployed infrastructure that aligned with the specific individuals and entities that were being targeted and launched complex social engineering schemes to lure targets to interact with the malicious sites. The lures that APT42 used were customized to include references to legitimate entities such as think tanks and, in at least one case, the threat actors referenced a specific target's name.

Conclusion

Attackers evolve, and so must defenses, but the fundamental principles that make up a robust security program remain critical. Some Iran-nexus threat actors continue to rely on credential harvesting and multifactor authentication (MFA) bypass for initial access. Any practice that raises the effort required to bypass MFA has a subsequent negative effect on threat actors. Enforcing phishing-resistant MFA methods, such as certificate-based authentication (CBA) and FIDO2 security keys, wherever possible, remains a core security practice—especially when it comes to privileged accounts. Similarly, as organizations continue to adopt cloud technology, a security-first design should be implemented to blend the business and operational needs with the security responsibilities of cloud operations. A design that seeks not only to define the security controls, but also ensures adequate visibility into all cloud-based activities, provides the necessary data for threat hunting, incident response, and ongoing monitoring. Finally, user awareness training—especially training that seeks to engender a community of responsibility when it comes to the security of the organization, its customers, and their data—is critical to the protection of any organization. Social engineering campaigns are becoming increasingly complex, and organizations should educate users on the ways in which they might be targeted outside of work-based perimeters.

As Iran-nexus threat actors continue to pursue cyber operations that align with the interests of the Iranian regime, they will alter their methodologies to adapt to the current security landscape. While evolutions in a threat actor's tactics, techniques, and procedures can result in temporary detection challenges, a comprehensive understanding of the factors that can fuel operations for these groups can help organizations in their threat hunting endeavors. Perhaps most importantly, collaboration across industries and sectors threatened by Iran-nexus actors is necessary to safeguard organizations from the risk posed by these groups.

Evolution of Data Theft in Cloud and Software-as-a-Service Environments

In recent years, Mandiant has observed a dramatic increase in cloud computing and software-as-a-service (SaaS) adoption, with organizations embracing these technologies for their scalability and flexibility. This shift, however, introduces a model where security responsibilities are shared between the provider and the customer in a highly contextual manner. While cloud and SaaS offerings bring numerous benefits, they also present unique security challenges for IT professionals, business leaders, and security practitioners tasked with securing these environments. Mandiant has observed attackers adapting to this shift in IT infrastructure and modifying the techniques they rely on for data theft. By understanding the evolving motivations and tactics, network defenders are able to embrace and build on practices that better address gaps in visibility, challenges with identity management, and complexities in strategic security plans.

Early Patterns of Data Theft

Data theft followed a relatively predictable pattern prior to the ready availability of cloud infrastructure. A typical scenario involved an attacker gaining access to a network, often through phishing or exploiting vulnerabilities in internet-facing systems. Once the attacker gained internal access to a targeted environment, they typically performed internal reconnaissance to map the network and identify valuable resources and data. Once they identified resources that fit their objective, threat actors escalated privileges to gain access to sensitive information stored within. That data would be copied to a compromised system in the environment and then stolen and stored on attacker infrastructure. This pattern was reliable enough to form the basis of the steps taken to identify threat actors during investigations.

To address the risk posed by threat actors, organizations relied on a combination of security controls and detection sources that they could build into their environment. At the time, this approach was made more effective by the existence of clear perimeters in a network and the relative simplicity of infrastructure that allowed it to be successfully managed by small organizations. Security instrumentation was developed over time to give greater degrees of visibility into network traffic, endpoint activity, and data transfers, while security information and event management (SIEM) platforms were developed to aggregate and highlight risk concerns. While this paradigm for detection and security served business needs well for decades, the value proposition of cloud-based technologies could not be ignored, and organizations leapt at the chance to do more for less. As client environments shifted away from traditional on-premises infrastructure to hybrid and cloud-native solutions, Mandiant observed threat actors shift their attack techniques in kind. While security fundamentals stayed relatively the same, many of the traditional security controls that were once effective in detection and mitigation of data theft started to fall behind.

Shifting Tactics: Attacker Adaptation and Exploitation

Throughout 2024, Mandiant observed attackers increasingly eschewing traditional on-premises network infiltration in favor of targeting cloud-based stores of centralized authority, such as single sign-on (SSO) web portals. When successful, these centralized authorities could grant a threat actor broad-scale access to an environment.
In the past, attackers would have to compromise a single system and move laterally through an environment before finally acquiring high-privilege access, such as domain admin credentials. The centralized nature of cloud identity and access management (IAM) technologies can provide a shortcut with fewer opportunities for exposure. High-value accounts can often be used to bridge access between cloud and on-premises environments. Attackers are targeting user credentials for cloud services and subsequently social engineering corporate help desk teams to reset passwords and enroll new multifactor authentication (MFA) devices to gain access to corporate identity solution portals. Threat actors such as UNC3944[41] used compromised SSO credentials to access virtual infrastructure management platforms and launched virtual machines (VMs) to support post-compromise activities and data theft. Mandiant has observed threat actors compromise on-premises accounts configured with certain cloud-related privileges and configurations.

In cases where the account is sufficiently privileged, these accounts can provide a very effective means to impact cloud environments from on-premises resources.

Where compromising a single privileged account can be a boon to threat actors, a threat actor gaining privileged access to SSO and identity management platforms can only be described as a windfall. These platforms are capable of granting broad-scale access across the cloud and SaaS environments with which they integrate. Once attackers gain access to these systems, they can often escalate privileges and pivot to other applications and services associated with these management consoles. Mandiant observed attackers with compromised SSO credentials add themselves to privileged groups that granted access to a wider range of SaaS applications.

Attackers are employing hybrid approaches, using both on-premises and cloud resources during their operations. During one investigation, Mandiant identified evidence that the threat actor discovered cloud access keys stored in plain text on the compromised on-premises network. The threat actor was able to use the keys to access and steal data from the client's cloud storage buckets. When the actor transferred the data they were stealing from the cloud buckets, they used a destination cloud bucket they controlled, which was hosted on the same platform. This helped the activity blend in with legitimate activity in the platform monitoring logs.

In addition to traditional social engineering of accounts with privileged access to on-premises solutions, Mandiant has observed a rise in the use of social engineering to target users that threat actors suspect have privileged access to SaaS environments. Deceiving a targeted user into providing credentials or approving MFA requests provides threat actors with an immediate escalation into cloud resources without having to compromise on-premises networks where security operations teams may have better visibility. This follow-on effect of targeting seeks to exploit potential gaps in understanding and visibility, while quickly accelerating the speed at which a threat actor can complete their mission objectives. The more a customer understands their subscription, the breakdown of responsibilities, and the means through which an investigation may be performed, the better prepared they are to not only withstand attacks, but to investigate them as well.

Managing Responsibilities and Risk

As the value presented by cloud infrastructure has become more apparent, situations have emerged where the priority of business operations has grown at a pace that outstrips the ability of security teams to identify risk and design security solutions. An area where this may become apparent is in the identification of realms of responsibility. Cloud platforms function under a shared responsibility model, where the responsibility for securing the environment stack is divided between the customer and the provider. The shared responsibility model, when not well understood, can lead to unmanaged risk and significant impacts to an investigation in the event of a compromise.

Not fully understanding the shared responsibility model may lead organizations to make assumptions that can damage their security posture. If an organization mistakenly believes that security is the sole responsibility of the provider, the security of the data, applications, and access controls in their environment can be placed at risk. This can be critical when the sensitivity of the data implies requirements under legal frameworks, such as the Sarbanes-Oxley Act. While the provider is commonly responsible for the security of the underlying infrastructure, customers are ultimately responsible for the security of their data and the applications they build. Fully understanding the organizations' responsibilities regarding security in a shared responsibility model is a critical aspect of designing secure environments.

In a similar vein, organizations should ensure they have a full accounting of where necessary log data is generated and by whom. It is an organization's responsibility to understand logging requirements from a forensic and regulatory perspective. It is important to collaborate with the cloud or SaaS provider to understand and verify the regulatory requirements with which they are compliant. Not all subscription levels provide the detail necessary to fully capture relevant information. Ensuring your subscription matches your requirements will assist with not only regulatory compliance, but security visibility.

Many legal frameworks related to the security of sensitive data have log generation and retention requirements. The quality and storage of cloud and SaaS-generated telemetry can also affect the pace of investigations into suspicious activity.

Mandiant has encountered multiple organizations that do not fully understand the implications of their specific subscription levels within the cloud and SaaS platforms they use. Even logs critical to audit logging for SaaS applications can be dependent on the customer's subscription level. Many investigations into cloud environments have been slowed or otherwise negatively impacted when the assumed logging level does not match the reality of the subscription service. Audit logging provides substantial value to network defenders tasked with monitoring for and investigating suspicious activity. The quality and quantity of the recorded logs can greatly decrease the time required to resolve an investigation and increase the confidence in the findings. The better a customer understands the features included at their subscription level, the associated breakdown of responsibilities, and how it may affect their visibility into critical areas of security, the better prepared they will be to identify risk and make informed changes.

Visibility Challenges in the Cloud

One of the most significant challenges in securing cloud environments is gaining the appropriate level of visibility into the environment. Where traditional environments have clear boundaries and choke points that could be instrumented, cloud environments scale more broadly and require an in-depth understanding of a variety of logging options. While the verbosity and availability of logs can vary greatly depending on the provider and the customer's subscription level, some log sources should be prioritized for collection and auditing.

Organizations should strive for logging that encompasses user logins and logouts, data access and modifications, administrative actions, system/configuration changes, and other security-related events. These logs should capture details such as timestamps, user identities, IP addresses, device information, and specific actions performed. Increased logging may lead to additional costs tied to a combination of cloud storage, processing, and service tiers. These costs have a ripple effect impacting managed service provider (MSP) pricing and organizations with limited security budgets.

Organizations should also take into consideration regulatory requirements and security capabilities when determining an appropriate log retention period. Ideally, customers should have easy and secure access to these logs with the ability to search, filter, and export data for analysis. The ability to integrate logs with SIEM tools for centralized log management, correlation, and analysis is also highly desirable as it supports not only business operations but investigative activities as well.

To monitor cloud environments effectively and detect potential data theft attempts, organizations should ensure comprehensive logging is enabled across their cloud services. The following log sources provide necessary visibility into various aspects of cloud infrastructure and should be enabled and regularly reviewed. Whenever possible, configuring alerts for suspicious log events and enabling timely detection and response to potential security incidents can help minimize the impact of successful attacks.

Network Traffic Logs

VPC Flow Logs

```
VPC flow logs (GCP and AWS)
NSG flow logs
VNet flow logs (Azure)
```

VPC flow logs capture information about IP traffic flowing to and from network interfaces within your virtual private cloud (VPC) or virtual network (VNet). They are essential for detecting unusual traffic patterns, identifying potential command-and-control (C2 or C&C) communication, and understanding network access to sensitive resources.

Verify that flow logs are enabled for each VPC, subnet, or network interface as needed.

Firewall Logs

Firewall logs, whether from a dedicated network firewall or integrated with your VPC/VNet, record details about traffic, which is allowed or denied based on your firewall rules. These logs help monitor network access to your resources and identify potential attempts to bypass security controls.

Verify that logs are stored in a centralized location, such as a SIEM.

Storage Access Logs

```
Cloud Storage access logs (GCP)
S3 Server access logs (AWS)
Storage Analytics logs (Azure)
```

Storage access logs provide detailed records of requests made to your cloud storage buckets. They are crucial for identifying unauthorized access, the scope of data exposure, and understanding how data is being accessed and used.

Confirm that access logging is enabled for all sensitive data storage buckets.

Compute and Resource Monitoring

```
gcloud logging API (GCP)
CloudWatch metrics (AWS)
Azure monitor metrics
```

While not traditional logs, these services provide performance and operational metrics for various cloud resources, including compute instances and storage volumes. Monitoring these metrics can help identify unusual resource utilization, which could indicate malicious activity, such as cryptomining or unauthorized data processing.

Confirm that logging is set appropriately, validated, and in a place where security personnel can review.

Audit Logging

```
Cloud Audit Logs (GCP)
CloudTrail Logs (AWS)
Azure Activity Logs
```

Audit logs record API calls and management actions made within a cloud environment and provide an audit trail of who did what and when. These logs are critical for detecting unauthorized configuration changes, privilege escalation attempts, and other suspicious administrative activity.

Ensure that the logs are activated for any and all cloud environments and in a place where security personnel can review.

Database Logs

Database logs can record accesses and commands executed against databases. Databases in both traditional and cloud technologies provide a target opportunity for sensitive information and are frequently an area containing visibility gaps.

Enable database-specific audit logs to monitor access and activity within your managed databases.

Identity and Access Management Logs

Many cloud providers offer specific logs related to IAM activities. IAM logs could include logs for authentication events, authorization failures, and changes to IAM policies.

Traditional Technique	New Cloud Adaptation
Internal Reconnaissance via SMB Scanning (MITRE T1135)	Cloud Storage Object Discovery (MITRE T1619)
Data Staging (MITRE T1074)	Modify Cloud Compute Infrastructure (MITRE T1578)
Data Collection from Local/ Network Systems (MITRE T1005/T1039)	Data collection directly from cloud storage services like S3, Azure Blob Storage, or Google Cloud Storage (MITRE T1530)
Exfiltration Over C2 Channels (MITRE T1041)	Data exfiltration directly to cloud storage services (MITRE T1567.002) or attacker-controlled accounts (MITRE T1537), blending exfiltration traffic with legitimate cloud usage

Table 4.1: Traditional data theft technique adaptations for cloud and SaaS environments

It is recommended to not only capture and store any IAM-specific logging but to set up relevant alerts and monitoring for security personnel to continuously review and audit.

Adapting Traditional Methods to the Cloud

While threat actors continue to evolve to meet new technologies, they are not abandoning their tried-and-true data theft techniques. Instead, they are simply adapting them to the cloud environment, creating a hybrid approach that leverages both on-premises and cloud resources.

Conclusion

The migration to cloud and SaaS environments has fundamentally changed the landscape of data theft. Attackers are adapting quickly, exploiting potential complexities of cloud infrastructure and security to their advantage. Relying solely on traditional security approaches designed for on-premises environments can lead organizations into areas of unsuspected risk. A security-first approach to cloud adoption is essential. By understanding the evolving threat landscape, implementing robust security controls, and fostering a culture of security awareness, organizations can reduce the risks of cloud data theft and harness the full potential of cloud computing.

Common Themes in Cloud Compromise Investigations

As organizations migrate to the cloud, protecting cloud and hybrid environments has grown increasingly complex. Organizations often look at their cloud infrastructure in isolation, focus on cloud-native controls, and aim to secure data and operations within the cloud itself. However, the evolving threat landscape is challenging the efficacy of this approach.

As a result, threat actors are capitalizing on misconfigurations that extend beyond the cloud's perimeter. By abusing these misconfigurations, attackers are able to gain access to cloud environments. This can be seen even in organizations with mature cloud security instrumentation. For example, Mandiant has encountered environments where the customer has deployed endpoint detection and response (EDR) tooling across all cloud-hosted virtual machines. With administrative access to the EDR managed through a federated identity provider, protections are often not designed to secure the EDR admin console in the event the identity store is compromised. Were an attacker to compromise the identity store in an environment such as this, they would be able to access the virtual machines (VMs) in the cloud through the EDR agents directly. This example, taken from frontline investigations performed by Mandiant, demonstrates how a compromise outside the boundary of the cloud environment can lead to a compromise of workloads in the cloud.

In 2024, Mandiant responded to more breaches that involved a cloud component than ever before. In the investigations Mandiant performed, three major themes contributed to threat actor successes in these environments:

1. Identity solutions that lack sufficient security controls
2. Improperly secured on-premises integrations
3. Poor visibility into extended cloud attack surface

Taken as a whole, these factors signal a need for a security approach that bridges the gaps between on-premises and cloud, while also recognizing that the cloud's attack surface is not isolated, but part of an interconnected ecosystem that demands proactive integrated defenses.

Securing Identities

Identity in cloud and/or hybrid environments serves as the first line of defense as many cloud incidents stem from compromised identities. Typically, these incidents originate from two key weaknesses: an identity architecture that does not protect against the use of compromised credentials and identity practices that include policies attackers can exploit.

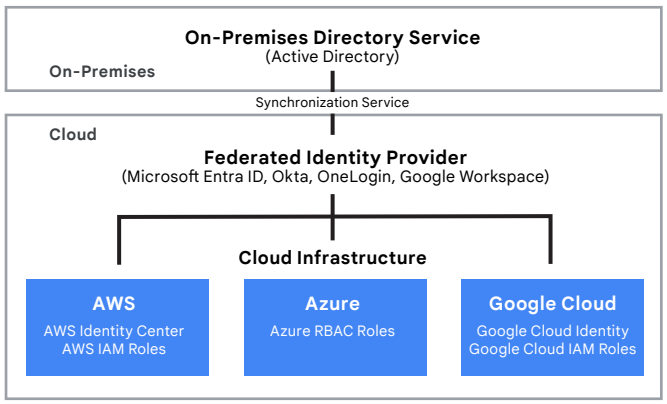

Figure 5.1: A common organizational identity architecture

Identity Architecture

A common organizational identity architecture typically includes an on-premises directory service, a federated identity provider, and a cloud Identity and Access Management (IAM) infrastructure service.

Organizations often adopt this architecture to unify their identity program and streamline authentication and authorization across all layers. While this setup is convenient, it can also introduce attack vectors that attackers frequently exploit. Attackers often target on-premises directory services, particularly when those services are used to manage and administer cloud environments. This creates a critical point of failure that can compromise

the entire system. Once an on-premises identity store is compromised, attackers can reuse those stolen credentials to access and compromise cloud resources directly.

Identity Practices

Attackers often seek the easiest and most efficient ways to compromise privileged identities and execute their attack chain—whether through malware deployment, data theft, or other malicious activities. The most common methods of identity compromise include brute forcing using common/guessable passwords, replaying stolen credentials from a previous breach, credential stuffing, phishing, and social engineering. Additionally, improperly secured identity practices often serve as a path of least resistance when attackers need to escalate privileges during a compromise. Mandiant categorizes commonly abused identity practices into three major areas: multifactor authentication (MFA), self-service, and third-party identities.

Mandiant regularly observes that organizations are not protecting privileged accounts with MFA. The absence of MFA leaves these accounts vulnerable to basic credential attacks, such as password spraying and credential stuffing. Even when implemented, MFA methods such as SMS, phone calls, or push notifications are susceptible to a variety of bypass techniques. These include adversary-in-the-middle (AiTM) attacks, account takeover via manipulation of the MFA registration process, social engineering, SIM swapping, intercepting MFA codes, and exploitating MFA fatigue. Additionally, many organizations do not secure the MFA registration and modification process sufficiently, which allows attackers in possession of compromised valid credentials to register their own MFA methods and continue operating undetected.

Mandiant has frequently observed attackers exploit password reset portals and related technologies to obtain credentials that grant them direct access to targeted organizations. Portals that are only protected by single-factor authentication or those that can be accessed from any device or location are particularly vulnerable to password-spraying attacks. Additionally, systems like interactive voice response (IVR), which rely on limited verification data such as date of birth, corporate information, employee IDs, or Social Security numbers, can be easily bypassed through social engineering campaigns.

Many organizations depend on third-party vendors, such as managed service providers (MSPs), to manage elements of their cloud environments. While external partners can streamline data, infrastructure, or security operations, granting them unlimited and unrestricted access often introduces considerable risk. Attackers frequently set their sights on third-party providers in the hopes that by compromising a single vendor, they can open pathways into multiple downstream organizations.

Organizations that lack sufficient controls around access to critical cloud data and infrastructure expose their identity stores to even greater risk. Because it is difficult to differentiate between compromised and legitimate credentials, security surrounding access should be commensurate with the sensitivity of the resources. By increasing the level of effort required to authenticate and interact with critical data and infrastructure, additional onus is applied to threat actors seeking to compromise the environment. Critical identity measures, such as privileged identity management (PIM) and phishing-resistant MFA, are relatively simple to implement and substantially improve security but require significant operational load to maintain and operate. Tying access to specific geographical locations or requiring privileged access workstations creates additional conditions that a threat actor must meet in order to gain access.

An aspect that sometimes gets overlooked is the security risk posed by members of the extended workforce. As organizations cannot enforce security controls on systems they do not own, the resources that contractors and vendors interact with should be tightly controlled. This includes enforcing limitations on the remote access management tools that are permitted to access critical resources and ensuring that a clear barrier between full-time employees and the extended workforce exists. A common way to accomplish this is to onboard third-party vendors into their own identity store separate from the corporate identity store.

On-Premises Integrations

As organizations deploy cloud infrastructure, it's common to create integrations with on-premises infrastructure to reduce friction for users and allow network and compute connectivity with existing systems. While this architecture has operational benefits, if an attacker is able to gain access to either of these environments, the integration could allow vertical movement between cloud and on-premises or vice versa. Mandiant has regularly observed evidence of threat actors having crossed the on-premises to cloud boundaries during intrusions. While threat group motivations may vary, the risk presented by not securing integrations has been demonstrated by prolific threat groups such as APT29, UNC3661, and UNC3944 crossing environments as they pursue

operational objectives. Even with state-of-the-art cloud security controls, improperly secured integrations with on-premises systems can allow an attacker to bypass these controls and compromise a cloud environment. These integrations can be broken down into two main categories: trusted service infrastructure and compute and network integrations.

Trusted Service Infrastructure

Trusted service infrastructure is typically associated with the management interfaces for platforms and technologies that provide core administrative services. Examples of trusted service infrastructure include:

- Asset and patch management tools
- Network management tools and devices
- Backup technologies
- Security tooling
- Virtualization consoles
- Privileged access management systems

As these are already associated with legitimate infrastructure within an environment, attackers will often target these platforms and abuse their intended functionality. Mandiant has observed attackers targeting trusted service infrastructure to pivot between cloud and on-premises infrastructure.

Compute and Network Integrations

Compute and network integrations are commonly used when organizations leverage infrastructure-as-a-service (IaaS) cloud components that are tightly integrated with on-premises environments. These integrations can allow an attacker that has compromised on-premises servers or virtual machines (VMs) to gain access to cloud VMs. Often in this scenario, the fact that these VMs are hosted in the cloud does not affect the attacker's techniques or motivations. For example, an attacker that has compromised an Active Directory privileged user account could impact a domain-joined VM hosted in the cloud. This could be via Group Policy Object deployment or, if the VMs share network connectivity, the attacker could remotely access the machine over RDP or SSH from the on-premises network.

Extended Cloud Attack Surface

Mandiant has often observed that organizations manage their attack surface from the perspective of a defined network boundary or perimeter. While network exposure remains a risk, the attack surface in cloud environments extends further.

The cloud attack surface encompasses the data attackers can enumerate about an organization's cloud environment. This includes details about identities, security configurations, settings, and resource configurations. This information is often accessible outside a network perimeter to low-privileged or even unauthenticated users. Freely available tools can collect significant volumes of data regarding cloud tenants if not properly secured.

Credential sprawl, including long-lived service account keys, also forms a critical component of the cloud attack surface. Inadvertent publishing of these credentials in public code repositories, shared documents, or other insecure locations often provides initial access and lateral movement opportunities. In addition, these credentials are often collected and posted for sale on dark web forums and chats. This is especially risky when cloud service accounts are assigned default or basic roles, such as Owner or Contributor. Organizations that do not centrally manage and secure service account credentials are susceptible to these types of attacks. Mandiant often encounters environments where service accounts are not properly documented and a baseline of their use does not exist. This can make recovery of a compromised environment a high-friction process as the ability to rotate credentials is slowed.

Lastly, publicly exposed and accessible resources expand the cloud attack surface. This can be from the perspective of both IaaS and platform-as-a-service (PaaS) components.[42] Risks with IaaS typically arise from VMs with public IP addresses and firewall rules allowing traffic from the internet on administrative ports. In PaaS environments, where the cloud provider manages the underlying infrastructure, misconfigured API or resource sharing can pose significant risks. These misconfigurations can allow access from external accounts or even anonymous access from the internet.

These factors require organizations to identify and reduce their cloud attack surface proactively and use tools that provide views into their environment similar to what an attacker would see. Cloud security posture management platforms have many valuable features, including the ability to provide a comprehensive inventory of cloud resources. This enables organizations to build a cloud asset management program, set standards on what should be exposed publicly, and then detect and remediate a non-compliant resource. Many platforms have attack surface management capabilities that provide visibility into internet-accessible resources, what software is running, and if there are vulnerabilities or entry points.

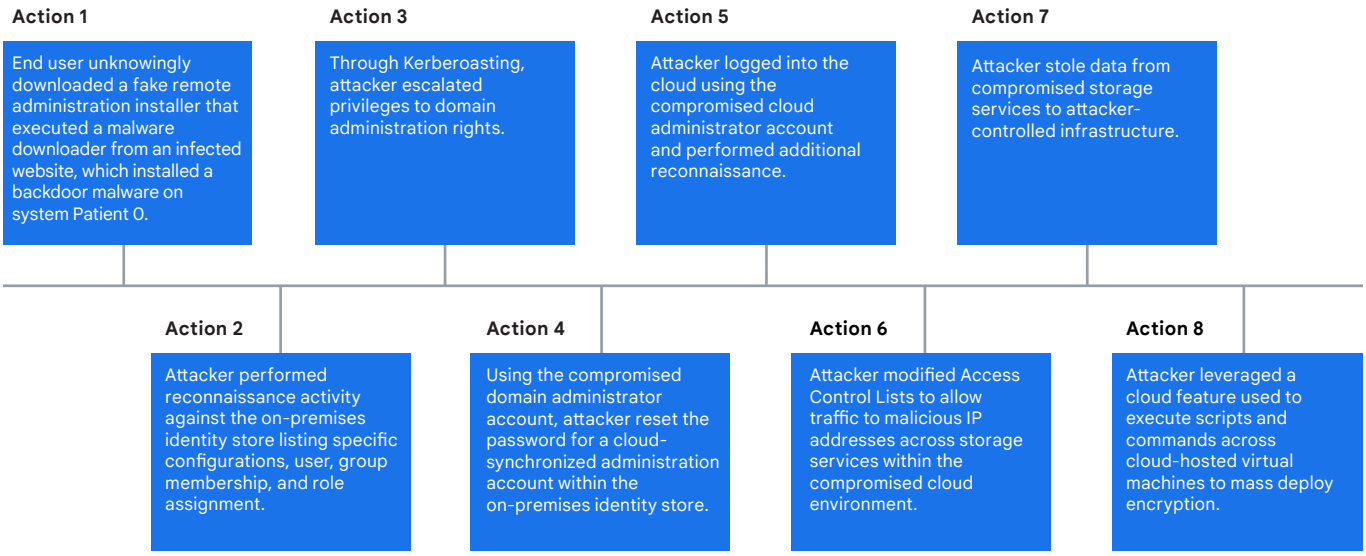

Figure 5.2: Incident Response Case Study

Incident Response Case Study

Mandiant was engaged by a customer to respond to an incident that included improper access to the customer's cloud environment. Mandiant incident responders identified evidence of a threat actor moving through the customer's hybrid environment and bypassing security controls to steal data before mass deploying ransomware.

The attack chain began when a user unknowingly downloaded a fake remote administration installer, resulting in the installation of a backdoor. The attacker then conducted reconnaissance of the on-premises identity store and used Kerberoasting to escalate privileges to Domain Administrator.

Due to the fact that the targeted organization leveraged their on-premises identity store to create and manage cloud administrator accounts, the attacker was able to acquire the password for a cloud administrative account. With this foothold established, the attacker moved into the cloud environment, performed additional reconnaissance, and modified access control lists to allow for communications to the attacker-controlled external infrastructure. This communication channel was then used for data theft. Lastly, the attacker leveraged a native cloud feature commonly used to deploy scripts to initiate a large-scale ransomware attack by encrypting cloud-based VMs.

Conclusion

Mitigating evolving attacker techniques in the cloud requires more than a single tool, configuration, or control. It demands a comprehensive, multilayered approach that includes carefully applied restrictions, hardening measures, ongoing detection strategies, and proactive response actions. By integrating these protections across every layer—from managed identities and resources to network and endpoint defenses—organizations can build a resilient security posture that anticipates and mitigates the complexities of today's hybrid threat landscape.

Security Recommendations for Diverse Cloud and Hybrid Environments

In any given year, Mandiant consultants respond to, assess, and advise thousands of clients across the various consulting services we provide. Engagements that include cloud or software-as-a-service (SaaS) components have become the norm as customers have expanded their environments to capitalize on the value presented by cloud technologies. While each environment is unique and poses its own challenges for security professionals, over the years, Mandiant consultants have identified a set of recommendations that can help provide a baseline for better security across diverse cloud and hybrid environments. By designing environments around identity and infrastructure controls that limit potential impacts of common threat actor activity, organizations can better secure their environments and meet their critical day-to-day business needs. Pairing controls with logging and detection capabilities that can provide substantial insight into activity in the environment gives network defenders the necessary visibility to validate controls, monitor for anomalies, and engage in threat hunting activities.

1. Reduce the Impact of Stolen Credentials

Stolen and compromised credentials are a common initial access vector used by threat actors regardless of skillset or objective. Organizations should train employees on the risks of password reuse and secure password management practices, which should be made available to employees and refreshed regularly. Technical controls applied to how users authenticate to and within the environment provide an additional layer of identity security. Organizations can reduce the impact of compromised credentials by implementing access control policies that evaluate necessary conditions dynamically before granting access for a user. Such conditions should include the following:

- All applications and resources should only be accessed through managed endpoints that are compliant with organizational policies.
- Modern authentication clients and protocols that require and enforce multifactor authentication (MFA) challenges should be configured for all accounts.
- Access to applications should require phishing-resistant MFA, such as FIDO2 security keys.
- Privileged account access should be restricted to known locations, IP addresses, and specific devices.
- All access requests should be evaluated for potential indicators of compromise prior to issuing authentication tokens.
- The lifespan of individual sessions should be limited to short periods of time and require re-authentication with MFA upon expiration.
- Cyber threat intelligence data, such as credential monitoring services, should be integrated into identity management to detect accounts that have been compromised and automatically expire active sessions.

Single-Factor Authentication (SFA) Permitted?	Weak Multifactor Authentication (MFA) Enforcement?	Strong MFA Enforcement?	Identity + Device Validation?	Multicontext Criteria Identity + Device + Geo + Origin Bound	Risk/Consequence of Stolen Credentials
Yes	No	No	No	No	High
Yes	No	No	Yes	No	Elevated
No	Yes	No	No	No	Elevated
No	Yes	Optional	No	No	Medium
No	No	Yes	No	No	Lower
No	No	Yes	Yes	No	Low
No	No	Yes	Yes	Yes	Lowest

Table 6.1: Control implementation to reduce the impact of stolen credentials

2. Protect Cloud Infrastructure from On-Premises Compromise

Organizations should focus on segmenting both cloud identities and resources to protect against on-premises compromise and vice versa. Privileged cloud accounts should not be synced to an on-premises identity store; instead, isolate privileged cloud accounts and limit their use to administrative tasks. Similarly, privileged accounts used to administer on-premises environments should not be synced to the cloud to protect against the same kind of lateral movement. Access to privileged accounts should be controlled using the principle of just-in-time access, which grants temporary privileges only when necessary.

For any trusted service infrastructure, the following actions are recommended:

- Limit the accounts that are allowed to authenticate to and access infrastructure tooling.
- Review and validate MFA enforcement for all accounts.
- Implement network restrictions to allow authentication and access from trusted IPs/networks only. Additionally, implement similar attribute-based access controls that can restrict access from specific identities, device types, and/or operating systems, where applicable.
- Create detections that focus on monitoring authentications and the activity performed within trusted service infrastructure.

3. Align Logging and Detection Strategy with Cloud Threats and Risks

Lack of visibility into cloud environments and logging limitations limit the efficacy of threat hunting, incident detection, and response activities. A comprehensive strategy defined to identify and address both general and specific risks to an organization's cloud infrastructure can help ensure investigations into suspicious activities are not impeded. Organizations should ensure that logging for storage bucket access, database access, and network flow logs are included in their configuration. Proactively reviewing the logging configurations beyond the default settings, validating the activity they capture, and centralizing logging to a SIEM should be a priority for security teams seeking better visibility into their cloud environment. Threat hunting and simulated attacker scenarios can help identify gaps in logging that may impede an investigation before they are able to negatively impact the process.

Threats to Web3 and Cryptocurrency

Malicious cyber operations involving Web3 technologies—cryptocurrencies, blockchains, and other decentralized user-centric technologies—are diverse, ranging from theft and money laundering to financing terrorist and military programs. Over the last three years, Mandiant has observed an increased targeting of the cryptocurrency industry, signaling an uptrend motivated by a variety of factors, such as its fast adoption, the security posture of the targets, and the inherent technical difficulty in disrupting these campaigns. While Web3 is not new, it is still considered a technology in emergence that is currently being integrated across industries beyond start-ups, including traditional finance institutions, the video game industry, and health and life insurance services.

Historically, emergent technology presents unique challenges to the entities adopting them, and Web3 is no exception. The financial sector is currently the most commonly targeted industry by threat activity and is also the largest adopter of Web3 technologies. Financial industries have introduced blockchain into their platforms,[43] created digital currencies,[44] and launched new products for financial markets[45] as financial regulations expand. This confluence of adoption and consistent threat actor activity has highlighted the need for additional scrutiny as organizations seek to protect their users, data, and digital assets.

The Rise of Web3: Opportunities and Risks

One of the inherent challenges organizations face when adopting new technologies is balancing the speed of integration while maintaining a robust security posture. As technologies mature, the methodologies through which threat actors target them also grow. Mandiant has observed threat actors targeting blockchain-native and blockchain-adopting industries in pursuit of a wide range of objectives. From leveraging cryptocurrency theft for financial gain[46] to the distribution of malicious code through censorship-resistant features of decentralized networks, threat actors are identifying ways to enrich themselves at the cost of the organizations and users exploring the benefits of Web3.

Cryptocurrency transactions, though recorded on a blockchain, pose challenges for both public institutions and private industry. Obfuscated fund flows and money laundering hinder large-scale tracing, while the immutability of smart contracts can prevent malicious code removal. These factors reduce threat actors' risk of identification, sanctions, or prosecution. Coupled with threat actors' adaptation to Web3, specialized phishing tools such as "drainers" targeting crypto wallets and Web3 projects create an ongoing threat, undermining trust in the ecosystem.

Democratic People's Republic of Korea Cyber Crime

In recent years, threat actors affiliated with the Democratic People's Republic of Korea (DPRK) have regularly targeted organizations and individuals who have adopted Web3 and cryptocurrency. In the past three years, Mandiant has investigated heists attributed to DPRK-nexus threat actors that resulted in over $500 million USD in stolen digital assets as a means of bypassing international sanctions. The focus on Web3 and cryptocurrency appears to be primarily financially motivated due to the heavy sanctions that have been placed on North Korea. Historically, the DPRK-nexus threat actor APT38 has been primarily responsible for attacks against financial institutions and some of the largest thefts of funds through cyberattacks.

DPRK-nexus threat actors appear to have access to a substantial cache of custom tooling written in a variety of languages, including Golang, C++, and Rust. These tools are often obfuscated with anti-analysis software, such as VMProtect and the open-source tool Garble. While obfuscating code is not a new tactic, nor is it impenetrable to analysis, raising the level of effort needed to reverse engineer their malware—or simply slowing the analysis—is sufficient cause for its use in many cases. Mandiant has also observed these threat actors deploying malicious tools designed for a variety of operating systems, including Windows, Linux, and macOS. Given the widespread use of macOS by developers, especially in the Web3 industry, the ability to compromise multiple

operating systems with custom tooling provides flexibility during cyber operations. These activities aim to generate financial gains, reportedly funding North Korea's weapons of mass destruction (WMD) program and other strategic assets.

UNC1069, UNC4899, and UNC5342 have adapted their methodologies to target members of the cryptocurrency and blockchain-development community more effectively. Specifically, they have come to target developers working individually and professionally on Web3-adjacent projects to obtain illicit access both to the cryptocurrency wallets of individual users and to the organizations that employ the impacted developers. UNC4736, on the other hand, is a prolific actor that targeted the blockchain industry in recent years by trojanizing trading software applications.

UNC1069

UNC1069, active since at least April 2018, targets diverse industries for financial gain. The group uses social engineering, often posing as investors from reputable firms on Telegram. UNC1069 has relied on spearphishing and social engineering to gain initial access and has been observed sending fake meeting invites (sometimes via compromised Telegram accounts) to Web3 and cryptocurrency organizations to gain illicit access to digital assets and cryptocurrency.

UNC4899

UNC4899, a suspected DPRK-nexus threat actor active since 2022, employs sophisticated social engineering and accesses via supply chain compromise. In 2024, UNC4899 targeted cryptocurrency professionals on social media with job postings for a prominent firm and gained access to Web3 organizations to steal digital assets. UNC4899 has previously conducted supply chain compromises to likely gain arbitrary access for financial gain.

UNC4736

UNC4736, a sophisticated North Korean threat actor, conducted a cascading software supply chain attack in 2022,[47] compromising a trading software entity and subsequently causing a second supply chain compromise that affected at least nine other organizations. This group has relied on trojanized trading and cryptocurrency software to gain network access for financial gain. UNC4736 also targeted decentralized finance platforms in 2024.

UNC5342

Mandiant began tracking UNC5342 in January 2024, following their social engineering campaign targeting software services, biotech, and media. UNC5342 distributed the BEAVERTAIL downloader via malicious cryptocurrency-themed NPM and Python packages hosted on GitHub. BEAVERTAIL downloads the INVISIBLEFERRET backdoor, granting UNC5342 extensive endpoint control.

Crypto Drainers and Smart-Contract Abuse

The new technologies and feature sets on which Web3 rely have created novel areas of expansion for threat actor techniques. Immutable elements of the blockchain and the decentralized nature of Web3 and cryptocurrency itself have been used by threat actors to create takedown-resistant infrastructure for use during campaigns. In traditional architectures, interorganizational cooperation is sufficient to perform takedown activities when a threat actor campaign is exposed. In 2024, the FBI and the US DOJ dismantled the 911 S5 botnet,[48] which affected millions of endpoints across the globe. However, by including components of Web3 technologies, such as smart contracts, threat actors can ensure that when a campaign is exposed, their activities can continue to operate as takedown activities coordinated over a decentralized ecosystem can be extremely difficult.

Smart contracts are an element that can be included in a blockchain to self-execute upon completion of a configured set of conditions that must be met. Once a smart contract is executed, the process is irreversibly recorded in the blockchain. Mandiant has observed threat actors using malicious smart contracts to steal digital assets and store key malware infrastructure elements within smart contracts.

Drainer operations often blend traditional tactics, such as phishing and social engineering, with malicious smart contracts that execute when a targeted user provides access to their cryptowallet. By luring users to approve malicious smart contracts, threat actors transfer the contents of a user's cryptowallet to one they control. While the Ethereum network has been the primary target of most drainer operations, Mandiant observed operations from DPRK-nexus threat actors expand their targeting of Ethereum to include the TRON and Solana blockchain platforms in 2023 and 2024, respectively. The DPRK-nexus threat actor UNC3782 commonly conducts large-scale phishing campaigns that focus on cryptocurrency industries. In 2023, UNC3782 conducted phishing operations against TRON users and transferred more than $137 million USD worth of assets in a single day. UNC3782 launched a campaign in 2024 to target Solana users and direct

them to pages that contained cryptocurrency drainers. Unlike their campaigns in 2023, however, Mandiant has not observed funds in Solana-based cryptocurrency wallets that are controlled by UNC3782, and the page hosting the Solana-based drainer was offline as of March 2024.

More than 1,200 fake sites associated with drainer operations have been created since January 2024. The financial gains found in drainer operations have led to the creation of drainer-as-a-service (DaaS) providers, who supply threat actors with the tools necessary to engage in drainer operations. DaaS providers advertise their services on underground forums and Telegram and receive a portion of the assets drained from user wallets as payment.

While the auto-executing nature of smart contracts can be used as a means to drain a user's assets in drainer operations, the immutability of smart contracts also allows threat actors to host takedown-resistant infrastructure on the blockchain. UNC5142, a financially motivated threat cluster tracked by Mandiant, targeted vulnerable WordPress websites[49] and injected code to retrieve data stored in a malicious smart contract. UNC5142's campaign ultimately resulted in the installation of infostealer malware and relied on the presence of the malicious smart contract, which contained second-stage code to fetch a payload from a remote server. This process of storing elements of an attack chain within smart contracts is commonly referred to as EtherHiding. When used during a campaign, EtherHiding allows for takedown-resistant infrastructure that can be updated as long as the smart contract is not executed and rendered immutable. Motivated threat actors leverage smart contracts to bypass traditional takedown measures and redirect their malware to use new infrastructure when existing infrastructure is disabled.

Conclusion

The rapid growth of blockchain technology across diverse industries has opened new avenues for adversaries to exploit. This includes targeting and manipulating the technology itself, while also enabling the misuse of cryptocurrencies. Ultimately, the nature of decentralized systems, coupled with a lack of security controls, lowers the perceived risks for malicious actors and poses challenges to law enforcement to intervene and react. Underground and darknet forums also play a role in criminalizing cryptocurrency by fueling an economy of illegal goods and services.

Crypto-native organizations in 2024 prioritized technical security for core wallet infrastructure and cryptographic controls, sometimes neglecting other standard controls. This focus, combined with rapid development and distributed workforces, often creates technical debt, expanding the attack surface. Challenges with evidence availability and quality, particularly regarding outsourced wallet infrastructure where log verification is often lacking, often hampers investigations. To address these issues, Mandiant recommends[50] enhanced transaction data monitoring and enrichment, combining with endpoint and security telemetry for better malicious activity detection.

Unsecured Data Repositories

While organizations pour resources into fortifying their perimeters against external threats, many overlook the basic security hygiene of their internal data repositories. These repositories often hold sensitive information, such as user credentials, financial data, and intellectual property, that are accessible to employees with standard privileges. This oversight creates an exploitable attack vector that enables threat actors to escalate privileges, steal data, and disrupt business operations.

Despite the risks posed by these caches of important data, this issue remains largely overlooked and is overshadowed by concerns of more sophisticated attacks. However, as organizations increasingly adopt cloud-based services and collaborative tools, the attack surface of unsecured repositories expands and further amplifies the risk. As threat actors target unsecured data repositories, a shift from a perimeter-centric security approach to a data-centric security approach has become necessary.

Mandiant uses similar tactics in Red Team engagements to model the common methodologies of threat actors. These engagements allow Mandiant to gain cross-industry security response data, and also bolster client threat defenses. The following red team and blue team case studies illustrate the efficacy of this attack vector on data repositories.

The Ripple Effect of Unsecured Data: A Red Team Case Study

Mandiant security assessments often identify sensitive data residing in readily accessible document repositories. Network file shares, SharePoint sites, Jira instances, Confluence spaces, and GitHub repositories often contain a wealth of valuable information (i.e., credentials, private keys, financial documents, personally identifiable information (PII), and intellectual property). This data, typically accessible to employees with standard privileges, presents a significant security risk that many organizations fail to recognize.

During Red Team engagements, Mandiant emulates advanced threat actors by performing custom and widely known tactics, techniques, and procedures (TTPs) in an attempt to compromise organizational data and achieve engagement objectives. Mandiant consultants take inspiration from observed threat actor activity to recreate the strategies that threat actors use during intrusions.

Mandiant was tasked by a customer to evaluate the security of a specific, cloud-native architecture backed by a massive data lake storing customer information. The customer detailed a set of objectives, which included successfully accessing specific data stores and compromising administrative systems. For the purpose of this project, Mandiant was provided with standard employee credentials that could be used to access the customer environment remotely.

In cloud-native environments, classic security assessment TTPs can be less effective as cloud-native architectures generally implement stronger authentication and authorization techniques than typical enterprise environments. Cloud environments that use a zero-trust model can also be more challenging to navigate for a threat actor as systems may be segregated into isolated virtual local area networks (VLANs). In addition, because cloud providers often maintain the underlying infrastructure and are more rigorous in their patching schedules, impactful vulnerabilities are often managed better within cloud-native environments than their on-premises counterparts.

As part of their reconnaissance efforts for the engagement, Mandiant generated a comprehensive list of document repositories that were accessible to the average employee. Mandiant identified a highly varied collection of data stores that allowed broad access, regardless of job role or group membership. Among the list, the customer's SharePoint document store and GitHub Enterprise repositories were prioritized for further analysis. SharePoint and GitHub are both widely adopted across many industries and organizations. Both platforms allow users to store arbitrary files, which are often used to support a diverse set of operations within an organization. Due to the often broad use case for these data stores, the data itself commonly outpaces the permissions that are applied to them. This can result in misconfigurations that allow for broader access than originally intended. Mandiant has observed that regular reviews of data stores, the

classifications of the data they contain, and subsequent review of access controls to match the classification are commonly overlooked steps in a security program.

Both SharePoint and GitHub provide built-in search functionality that allows for fine-grained querying, including the ability to filter keywords and file types. Mandiant incident responders regularly identify threat actors querying data repositories for file types likely to aid the attacker in various stages of the targeted attack lifecycle. Files ending in .pem or .key commonly store private keys that threat actors can use to access remote systems. Queries for filenames matching common Secure Shell (SSH) private key filenames, or even simple queries for files that include the word "password," have often been identified in browser search histories and application logs during an investigation.

A series of searches through the identified repositories led Mandiant to SSH private keys, application secrets, and user passwords that were stored in plain text with minimal access controls. By combining the stored secrets identified in GitHub with the SSH private keys from SharePoint, Mandiant was able to initiate a chain of lateral movement through the customer environment. With each system that Mandiant moved to, further searches and reconnaissance through technical documentation, password vaults, and credentials stored within a variety of virtual machines allowed Mandiant to progress one system closer to the objectives laid out by the customer. Network documentation identified in SharePoint was used as a navigation plane through which Mandiant tested credentials pulled from corporate password vaults. Identities that were compromised during each lateral move were subsequently used to access different sets of document stores, which fed into the cyclical process of identification, testing, and actioning. Mandiant continued to compromise systems within the environment until user credentials yielded access to privileged credentials, which then led to administrative credentials. Ultimately, Mandiant was able to engineer a path to the mission objectives by escalating their privileges to an administrator level and gaining access to sensitive data stores. Mandiant performed this attack chain without the use of malware, zero-day vulnerabilities, or any other more advanced attacker methodologies.

Exploitation in the Wild: An Incident Responder's View

Due to the nature of the information commonly stored in centralized data repositories, these repositories often present a valuable target to motivated threat actors, regardless of the mission objective. Mandiant tracks the motivations of thousands of threat actors identified through incident response engagements and intelligence-gathering operations. Mandiant has observed financially motivated threat actors, such as FIN11, UNC2891, and UNC3944, steal data from unsecured data repositories on which they can build high-price extortion demands from targeted organizations. On the other end of the spectrum, Mandiant has observed advanced persistent threat (APT) groups such as APT29, a threat actor attributed to Russia's Foreign Intelligence Service, steal data from information stores in pursuit of espionage objectives.

While unsecured data repositories are often overlooked by security teams, threat actors have recognized their inherent value as both a potential windfall for operational objectives and a cache of intel on their target environment. As such, threat actors across all levels of sophistication are likely to find factors that can drive the success of their operations. Depending on organizational needs, centralized data repositories may house critical information pertaining to day-to-day business operations. Data repositories often grow over time to contain more and more sensitive data as the use of the repository outstrips the original data classification against which the access controls were defined. Similarly, data lifecycle management processes are often deprioritized by operational teams, which results in organizations keeping data long past the business case under which the repository was originally established. Insufficiently safeguarding repositories that contain sensitive data inherently lowers the level of effort required for threat actors to pursue their mission objectives.

Financially motivated threat groups have historically relied on disruption of service through ransomware as a means to apply pressure to targeted organizations, with offers of relief available after a substantial payout. Over the years, organizations recognized the threat to their continued operation presented by ransomware and invested in technology such as early warning systems and disaster recovery to ensure they could return to an active state without an exorbitant payout. More recently, threat actors have responded in kind by adding a more material extortion to the mix. Instead of smash-and-grab ransomware operations where the time to encryption was made a priority, threat groups such as FIN11, UNC2891, and UNC3944 have progressed and prioritized selective data theft prior to encryption. The release of sensitive data stolen from the targeted organization is then used as additional leverage during the negotiation

for the decryption of the environment. Mandiant has even observed threat actors such as UNC3944 request follow-on payments to ensure stolen data is not leaked once payment has been made to decrypt the environment.

However, the direct impact of stolen data is not always as obvious when a threat actor is concerned more with espionage than with financial gain. APT29 is highly sophisticated and known for persistence in maintaining access to compromised environments, even after activity has been identified. As their remit is commonly more targeted toward the acquisition of valuable intelligence, targeting unsecured repositories that may contain sensitive data is a natural step for APT29, which can provide an abbreviated path toward their mission objectives. Mandiant has observed APT29 steal data on targeted personnel and critical infrastructure from data repositories where the level of protection did not match the classification of the data. Even in cases where the stolen dataset does not meet a threat actor's objectives, Mandiant has observed threat actors pursue data that simply aligns with the progress of their intrusion.

The ongoing support requirements of an organization's business and operational needs often rely on the quality and quantity of a set of well-maintained documentation. From network diagrams and troubleshooting guides, to full incident playbooks and application design documents, these stores of information, by necessity, exist to ensure the organization runs smoothly. Unfortunately, they also represent an added risk of information exposure, which threat actors rely on during the targeted attack lifecycle. Where manual reconnaissance of an environment can be noisy and risk exposing a threat actor's activity, manually reviewing network architecture documentation can provide the same if not better information to a threat actor. UNC2891, known for targeting environments with Linux and Solaris systems, has been observed using data from unsecured repositories to inform lateral movement in targeted networks. Similarly, Mandiant has observed APT29 steal information on systems of interest prior to moving laterally and compromising them.

Shifting to a Data-Centric Approach: Defensive Strategies and Recommendations

A review of recent security assessments performed by Mandiant revealed that roughly 46% of the engagements identified insecure storage of credentials or secrets as a risk factor. Given the diverse nature of the environments into which Mandiant is contracted and the breadth of security models encountered, findings regarding basic security hygiene tend not to cluster so dramatically.

While the traditional model of looking at a corporate computer network from the perspective of its perimeter and component systems has been valuable, augmenting that view with a layer for data residency and controls can help build more robust models. Focusing on where data resides—whether on-premises, in the cloud, or in separate software-as-a-service applications—should be a focal point for security teams. Given the data-centric nature of modern organizations, this task is not a trivial one and involves a multitiered approach:

1. **Perform inventory of data repositories:** Begin by pinpointing where sensitive data resides. This includes PII, financial records, corporate secrets, IT data, intellectual property, and anything subject to regulatory compliance (GDPR, HIPAA, etc.).
2. **Routinely audit data repositories:** Data repositories should be routinely audited with automated tools to identify exposed credentials and secrets.
3. **Implement robust access controls:** Blanket permissions are a commonly abused vector through which threat actors gain access to sensitive data. Ensure users have only the minimum accesses to data necessary to perform their jobs. Similarly, distinguish between users who require read access and those who require read/write access.
4. **Educate users:** Educate users about data security best practices, the importance of protecting sensitive data, and the consequences of data breaches. Employees should be trained to identify and report instances of sensitive data in open data repositories or secrets stored in code bases.
5. **Validate data is encrypted:** Encrypt data both in transit using protocols such as TLS/SSL and at rest to limit windows of exposure.
6. **Configure multifactor authentication (MFA):** Enforcing MFA for all accesses to sensitive data repositories adds an extra layer of security beyond passwords. For single sign-on (SSO)-based data repositories, ensure MFA is enforced when accessing SSO resources.
7. **Implement data loss prevention (DLP):** Consider implementing DLP solutions to prevent sensitive data from leaving your environment through observable network sessions such as email attachments or file transfers.

8. **Regularly audit the content of data repositories:** Review data repositories to ensure their contents match the classification of the original use case. Any sensitive data identified outside of secured containers should be logged, removed, and necessitate the start of a full search for similar data outside of secured locations. Data that is no longer needed for business purposes should be removed to keep the organization's data footprint within a manageable size. Automated tools can be used to facilitate data footprint reduction and track down sensitive files of interest.
9. **Implement zero trust and microsegmentation:** By adapting a zero-trust model in addition to microsegmentation, leaked credentials become significantly harder to abuse. Internal firewalls, context-aware access, and zero trust-based authentication all act as methods to restrict connectivity to a resource even if valid credentials are obtained.
10. **Perform dynamic secret management:** Tools that provision just-in-time access to secrets along with automated rotation after use, reduce the opportunities for a threat actor to misuse stolen credentials. Even if a credential is leaked, a dynamic secret management system should limit the impact by automatically rotating the credential and expiring active sessions. This technique also reduces administrative overhead along with providing detailed tracking of credential usage.
11. **Integrate CI/CD pipelines with dynamic secret management systems:** As software and systems are built and maintained, integrating continuous integration and continuous delivery/deployment (CI/CD) pipelines with dynamic secret management systems provides an opportunity to rotate credentials as infrastructure and assets move from development to production. Credentials for active, live systems with production data could be automatically rotated as code and configurations change, lessening the chance that a leaked credential remains valid.
12. **Perform regular security assessments:** Recurrent and regular security assessments help determine the impact and overall exploitability of any identified credentials. These tests also highlight how closely an organization follows their intended process, gaining a ground truth assessment of security control's effectiveness.

Conclusion

The presence of sensitive data within unsecured document repositories is pervasive and represents a significant yet often overlooked security risk. Despite investing heavily in perimeter defenses, improper controls applied to internal data stores can leave organizations and their data vulnerable to exploitation. Addressing and subsequently maintaining solutions to vulnerabilities in an environment helps reduce an organization's exposure to risk and provides a firm baseline of security on which further advancements can be built.

Expanding an organization's security measures to include the data that drives its success strengthens its existing security systems. While the perimeter of an environment may often represent the first chance to detect and inhibit threat actor activity, the systems that collect an organization's data can represent the last opportunity defenders have to prevent data theft. Placing barriers between threat actors or insider threats and sensitive data managed by an organization not only limits access but adds opportunities through which security teams can detect misuse. Comprehensive access controls, continuous monitoring, data tagging, and regular audits of repositories should form the starting point and inform the growth of a data-centric environment. By prioritizing data security across all platforms and cultivating a security-conscious culture, organizations can strengthen their overall security posture and better safeguard their valuable assets.

Conclusion

In 2024, we saw attackers take advantage of opportunities. This includes leveraging credentials obtained in infostealer campaigns for initial access, taking advantage of misconfigurations and weakly secured identities in hybrid environments, gaining access to data as a result of poor basic security hygiene, and targeting cryptocurrency and Web3 amidst its rapid adoption.

We also saw threat actors create opportunities, as seen with the Democratic People's Republic of Korea IT workers. These actors are brazen in their approach, notably targeting gaps in onboarding processes to obtain employment through deceptive means, and ultimately achieving their goals of funding the regime while also maintaining insider access to an organization.

Defending against the threats covered in M-Trends 2025 is no easy task. Effective cyber defense requires an extensive and multi-layered approach. Security teams must be rigorously tested through red team exercises and other simulations. Security teams should partner with Communications, Legal, and other relevant teams to conduct regular tabletop exercises to validate and improve incident response plans throughout the year. A cyber incident response retainer ensures immediate access to expert help, minimizing downtime and damage during a critical cyberattack.

Exploits (33%), stolen credentials (16%), and phishing (14%) were the most common initial infection vectors in our 2024 investigations. Foundational security practices, such as vulnerability management, least privilege, and system hardening, are essential. Organizations should build a comprehensive security program that covers all aspects of the enterprise, from cloud and on-premises environments to IT/OT systems and all assets, that is powered by strong detection and proactive threat hunting capabilities, and informed by impactful threat intelligence. And of course, employee education is a must.

The Mandiant mission is to help keep every organization secure from cyber threats and confident in their readiness. Our annual M-Trends report, featuring data and learnings from our engagements, plays a big part in advancing that mission. We will continue to share our frontline knowledge in M-Trends to improve our collective security awareness, understanding, and capabilities.

Mandiant, part of Google Cloud, has been at the forefront of cyber security and threat intelligence since 2004. Our incident responders are on the frontlines of the world's most complex breaches. We have a deep understanding of both existing and emerging threat actors, as well as their rapidly changing tactics, techniques, and procedures. Mandiant helps organizations quickly get back to business after a security breach and applies frontline expertise to guide effective threat detection, preparation, and to reduce business risk and build overall resiliency—before, during, and after an incident. Since 2010, Mandiant has been dedicated to publishing comprehensive trends based on our incident response engagements, providing critical insights into the evolving threat landscape through the M-Trends report.

If your organization suspects a cyber incident, or you are experiencing a security breach, please contact Mandiant for Incident Response Assistance.

MITRE ATT&CK

Techniques Related to Mandiant Targeted Attack Lifecycle, 2024

Mandiant's Targeted Attack Lifecycle is the predictable sequence of events cyber attackers use to carry out their attacks.

☁ indicates techniques in the Cloud matrix, introduced in ATT&CK v16.

Initial Reconnaissance

Reconnaissance

T1598: Phishing for Information ☁	1.3%		
T1595: Active Scanning ☁	0.6%	T1595.002: Vulnerability Scanning ☁	0.6%

Resource Development

T1588: Obtain Capabilities ☁	15.4%	T1588.003: Code Signing Certificates ☁	14.8%
		T1588.004: Digital Certificates ☁	0.4%
		T1588.007: Artificial Intelligence ☁	0.2%
T1608: Stage Capabilities ☁	12.3%	T1608.005: Link Target ☁	3.8%
		T1608.003: Install Digital Certificate ☁	3.2%
		T1608.001: Upload Malware ☁	1.7%
		T1608.006: SEO Poisoning	1.3%
		T1608.002: Upload Tool ☁	1.3%
		T1608.004: Drive-by Target ☁	0.6%
T1584: Compromise Infrastructure ☁	4.4%		
T1583: Acquire Infrastructure ☁	4.0%	T1583.003: Virtual Private Server ☁	4.0%
T1587: Develop Capabilities ☁	0.2%	T1587.003: Digital Certificates ☁	0.2%
T1585: Establish Accounts ☁	0.2%	T1585.002: Email Accounts ☁	0.2%

Initial Compromise

Initial Access

T1190: Exploit Public-Facing Application ☁	23.5%		
T1133: External Remote Services ☁	20.9%		
T1078: Valid Accounts ☁	19.5%	T1078.004: Cloud Accounts ☁	12.1%%
T1566: Phishing ☁	12.3%	T1566.002: Spearphishing Link ☁	5.5%
		T1566.001: Spearphishing Attachment	1.7%
		T1566.004: Spearphishing Voice ☁	1.5%
		T1566.003: Spearphishing via Service	1.3%
T1189: Drive-by Compromise ☁	4.4%		
T1199: Trusted Relationship ☁	0.8%		
T1091: Replication Through Removable Media	0.4%		
T1195: Supply Chain Compromise	0.2%	T1195.002: Compromise Software Supply Chain	0.2%
T1200: Hardware Additions	0.2%		

Establish Foothold

Persistence

T1133: External Remote Services ☁	20.9%		
T1078: Valid Accounts ☁	19.5%	T1078.004: Cloud Accounts ☁	12.1%
T1543: Create or Modify System Process	19.2%	T1543.003: Windows Service	11.0%
		T1543.004: Launch Daemon	0.4%
		T1543.002: Systemd Service	0.2%
T1098: Account Manipulation ☁	18.6%	T1098.007: Additional Local or Domain Groups	6.3%
		T1098.005: Device Registration ☁	4.7%
		T1098.004: SSH Authorized Keys	1.5%
		T1098.001: Additional Cloud Credentials ☁	0.2%
		T1098.003: Additional Cloud Roles ☁	0.2%
		T1098.006: Additional Container Cluster Roles ☁	0.2%
T1053: Scheduled Task/Job ☁	13.5%	T1053.005: Scheduled Task	12.7%
		T1053.003: Cron	0.8%
T1547: Boot or Logon Autostart Execution	11.0%	T1547.001: Registry Run Keys / Startup Folder	10.8%
		T1547.005: Security Support Provider	0.8%
		T1547.009: Shortcut Modification	0.6%
		T1547.002: Authentication Package	0.2%
T1505: Server Software Component	7.0%	T1505.003: Web Shell	7.0%
		T1505.004: IIS Components	0.2%
T1136: Create Account ☁	6.6%	T1136.001: Local Account	4.4%
		T1136.002: Domain Account	0.2%
T1574: Hijack Execution Flow	6.3%	T1574.011 Services Registry Permissions Weakness	5.5%
		T1574.002: DLL Side-Loading	0.8%
		T1574.001: DLL Search Order Hijacking	0.2%
T1546: Event Triggered Execution	3.6%	T1546.003: WMI Event Subscription	2.5%
		T1546.008: Accessibility Features	0.2%
		T1546.004: Unix Shell Configuration Modification	0.2%
		T1546.015: Component Object Model Hijacking	0.2%
		T1546.012: Image File Execution Options Injection	0.2%
T1556: Modify Authentication Proces	2.1%	T1556.006: Multi-Factor Authentication	1.1%
		T1556.009: Conditional Access Policies ☁	0.4%
T1037: Boot or Logon Initialization Scripts	0.8%	T1037.001: Logon Script (Windows)	0.2%
T1554: Compromise Client Software Binary	0.4%		
T1137: Office Application Startup ☁	0.2%	T1137.006: Add-ins ☁	0.2%

Escalate Privileges

Privilege Escalation

T1078: Valid Accounts ☁	19.5%	T1078.004: Cloud Accounts ☁	12.1%
T1543: Create or Modify System Process	19.2%	T1543.003: Windows Service	11.0%
		T1543.004: Launch Daemon	0.4%
		T1543.002: Systemd Service	0.2%
T1098: Account Manipulation	18.6%	T1098.007: Additional Local or Domain Groups	6.3%
		T1098.006: Additional Container Cluster Roles ☁	0.2%
T1055: Process Injection	15.0%	T1055.001: Dynamic-link Library Injection	0.6%
		T1055.003: Thread Execution Hijacking	0.6%
		T1055.012: Process Hollowing	0.4%
		T1055.004: Asynchronous Procedure Call	0.2%
		T1055.009: Proc Memory	0.2%
		T1055.002: Portable Executable Injection	0.2%
T1053: Scheduled Task/Job ☁	13.5%	T1053.005: Scheduled Task	12.7%
		T1053.003: Cron	0.8%
T1547: Boot or Logon Autostart Execution	11.0%	T1547.001: Registry Run Keys / Startup Folder	10.8%
		T1547.005: Security Support Provider	0.8%
		T1547.009: Shortcut Modification	0.6%
		T1547.002: Authentication Package	0.2%
T1134: Access Token Manipulation	7.6%	T1134.001: Token Impersonation/Theft	2.7%
T1574: Hijack Execution Flow	6.3%	T1574.011: Services Registry Permissions Weakness	5.5%
		T1574.002: DLL Side-Loading	0.8%
		T1574.001: DLL Search Order Hijacking	0.2%
T1546: Event Triggered Execution	3.6%	T1546.003: WMI Event Subscription	2.5%
		T1546.004: Unix Shell Configuration Modification	0.2%
		T1546.015: Component Object Model Hijacking	0.2%
		T1546.012: Image File Execution Options Injection	0.2%
		T1546.008: Accessibility Features	0.2%
T1037: Boot or Logon Initialization Scripts	0.8%	T1037.001: Logon Script (Windows)	0.2%
T1484: Domain Policy Modification ☁	0.8%	T1484.001: Group Policy Modification	0.8%
T1068: Exploitation for Privilege Escalation ☁	0.4%		
T1548: Abuse Elevation Control Mechanism	0.2%	T1548.002: Bypass User Account Control	0.2%

Internal Reconnaissance

Discovery

T1083: File and Directory Discovery	32.1%		
T1082: System Information Discovery ☁	24.5%		
T1033: System Owner/User Discovery	20.3%		
T1087: Account Discovery ☁	18.2%	T1087.002: Domain Account	8.2%
		T1087.001: Local Account	7.2%
		T1087.004: Cloud Account ☁	0.4%
T1016: System Network Configuration Discovery	17.5%	T1016.001: Internet Connection Discovery	9.9%
T1518: Software Discovery ☁	16.7%	T1518.001: Security Software Discovery ☁	0.6%
T1057: Process Discovery	16.7%		
T1012: Query Registry	15.0%		
T1622: Debugger Evasion	10.8%		
T1614: System Location Discovery	9.5%	T1614.001: System Language Discovery	4.9%
T1069: Permission Groups Discovery ☁	9.1%	T1069.002: Domain Groups	6.1%
		T1069.001: Local Groups	1.3%
		T1069.003: Cloud Groups ☁	0.6%
T1497: Virtualization/Sandbox Evasion	9.1%	T1497.001: System Checks	7.2%
T1482: Domain Trust Discovery	8.0%		
T1049: System Network Connections Discovery ☁	6.1%		
T1010: Application Window Discovery	5.5%		
T1007: System Service Discovery	5.3%		
T1018: Remote System Discovery	4.9%		
T1135: Network Share Discovery	3.8%		
T1046: Network Service Discovery ☁	2.3%		
T1124: System Time Discovery	1.3%		
T1580: Cloud Infrastructure Discovery ☁	1.1%		
T1619: Cloud Storage Object Discovery	0.8%		
T1615: Group Policy Discovery	0.6%		
T1538: Cloud Service Dashboard ☁	0.6%		
T1654: Log Enumeration	0.4%		
T1217: Browser Bookmark Discovery	0.2%		
T1201: Password Policy Discovery	0.2%		
T1120: Peripheral Device Discovery	0.2%		
T1613: Container and Resource Discovery ☁	0.2%		
T1040: Network Sniffing	0.2%		
T1652: Device Driver Discovery	0.2%		

Lateral Movement

Lateral Movement

T1021: Remote Services	33.2%	T1021.002: SMB/Windows Admin Shares	21.8%
		T1021.001: Remote Desktop Protocol	21.1%
		T1021.004: SSH	12.3%
		T1021.006: Windows Remote Management	1.3%
		T1021.005: VNC	1.1%
T1570: Lateral Tool Transfer	1.3%		
T1550: Use Alternate Authentication Material ☁	1.3%	T1550.002: Pass the Hash	1.1%
		T1550.001: Application Access Token ☁	0.2%
T1534: Internal Spearphishing ☁	0.6%		
T1072: Software Deployment Tools	0.4%		
T1091: Replication Through Removable Media	0.4%		
T1210: Exploitation of Remote Services	0.2%		

Maintain Presence

Persistence

T1133: External Remote Services ☁	20.9%		
T1078: Valid Accounts ☁	19.5%	T1078.004: Cloud Accounts ☁	12.1%
T1543: Create or Modify System Process	19.2%	T1543.003: Windows Service	11.0%
		T1543.004: Launch Daemon	0.4%
		T1543.002: Systemd Service	0.2%
T1098: Account Manipulation ☁	18.6%	T1098.007: Additional Local or Domain Groups	6.3%
		T1098.005: Device Registration ☁	4.7%
		T1098.004: SSH Authorized Keys	1.5%
		T1098.001: Additional Cloud Credentials ☁	0.2%
		T1098.003: Additional Cloud Roles ☁	0.2%
		T1098.006: Additional Container Cluster Roles ☁	0.2%
T1053: Scheduled Task/Job ☁	13.5%	T1053.005: Scheduled Task	12.7%
		T1053.003: Cron	0.8%
T1547: Boot or Logon Autostart Execution	11.0%	T1547.001: Registry Run Keys/Startup Folder	10.8%
		T1547.005: Security Support Provider	0.8%
		T1547.009: Shortcut Modification	0.6%
		T1547.002: Authentication Package	0.2%
T1505: Server Software Component	7.0%	T1505.003: Web Shell	7.0%
		T1505.004: IIS Component	0.2%
T1136: Create Account ☁	6.6%	T1136.001: Local Account	4.4%
		T1136.002: Domain Account	0.2%
T1574: Hijack Execution Flow	6.3%	T1574.011: Services Registry Permissions Weakness	5.5%
		T1574.002: DLL Side-Loading	0.8%
		T1574.001: DLL Search Order Hijacking	0.2%
T1546: Event Triggered Execution	3.6%	T1546.003: WMI Event Subscription	2.5%
		T1546.008: Accessibility Features	0.2%
		T1546.004: Unix Shell Configuration Modification	0.2%
		T1546.015: Component Object Model Hijacking	0.2%
		T1546.012: Image File Execution Options Injection	0.2%
T1556: Modify Authentication Process	2.1%	T1556.006: Multi-Factor Authentication ☁	1.1%
		T1556.009: Conditional Access Policies ☁	0.4%
T1037: Boot or Logon Initialization Scripts	0.8%	T1037.001: Logon Script (Windows)	0.2%
T1554: Compromise Client Software Binary	0.4%		
T1137: Office Application Startup ☁	0.2%	T1137.006: Add-ins ☁	0.2%

Mission Completion

Collection

T1560: Archive Collected Data	12.9%	T1560.001: Archive via Utility	6.3%
		T1560.002: Archive via Library	0.8%
T1213: Data from Information Repositories ☁	12.3%	T1213.002: Sharepoint ☁	7.6%
		T1213.003: Code Repositories ☁	0.4%
		T1213.001: Confluence ☁	0.2%
T1114: Email Collection ☁	7.4%	T1114.002: Remote Email Collection ☁	0.6%
		T1114.003: Email Forwarding Rule ☁	0.6%
		T1114.001: Local Email Collection	0.2%
T1074: Data Staged ☁	6.1%	T1074.001: Local Data Staging	5.3%
		T1074.002: Remote Data Staging ☁	0.6%
T1056: Input Capture ☁	3.8%	T1056.001: Keylogging ☁	3.8%
T1113: Screen Capture	3.4%		
T1039: Data from Network Shared Drive	2.3%		
T1115: Clipboard Data	2.1%		
T1005: Data from Local System	1.5%		
T1125: Video Capture	1.5%		
T1602: Data from Configuration Repository ☁	1.1%	T1602.001: SNMP (MIB Dump) ☁	0.2%
		T1602.002: Network Device Configuration Dump ☁	1.1%
T1530: Data from Cloud Storage ☁	0.6%		
T1123: Audio Capture	0.6%		
T1119: Automated Collection	0.2%		

Mission Completion

Exfiltration

T1567: Exfiltration Over Web Service	2.7%	T1567.002: Exfiltration to Cloud Storage	1.5%
		T1567.001: Exfiltration to Code Repository	0.4%
T1041: Exfiltration Over C2 Channel	1.1%		
T1020: Automated Exfiltration ☁	0.4%		

Impact

T1486: Data Encrypted for Impact ☁	24.1%		
T1657: Financial Theft	10.8%		
T1489: Service Stop	8.5%		
T1529: System Shutdown/Reboot	5.9%		
T1490: Inhibit System Recovery	4.9%		
T1565: Data Manipulation	4.4%	T1565.001: Stored Data Manipulation	4.4%
T1485: Data Destruction ☁	3.2%		
T1496: Resource Hijacking ☁	3.0%		
TT1491: Defacement ☁	1.3%	T1491.002: External Defacement ☁	0.8%

Other

Command and Control

T1071: Application Layer Protocol	21.4%	T1071.001: Web Protocols	16.9%
		T1071.004: DNS	4.9%
T1105: Ingress Tool Transfer	20.9%		
T1095: Non-Application Layer Protocol ☁	18.0%		
T1572: Protocol Tunneling	8.7%		
T1573: Encrypted Channel	5.9%	T1573.002: Asymmetric Cryptography	5.7%
		T1573.001: Symmetric Cryptography	0.2%
T1090: Proxy ☁	3.6%	T1090.003: Multi-hop Proxy ☁	1.5%
		T1090.001: Internal Proxy	0.6%
T1219: Remote Access Software	2.3%		
T1102: Web Service	1.3%	T1102.002: Bidirectional Communication	0.2%
T1132: Data Encoding	0.8%	T1132.001: Standard Encoding	0.8%
T1571: Non-Standard Port	0.6%		
T1008: Fallback Channels	0.2%		
T1104: Multi-Stage Channels	0.2%		

Execution

T1059: Command and Scripting Interpreter ☁	42.7%	T1059.001: PowerShell	24.5%
		T1059.003: Windows Command Shell	15.0%
		T1059.004: Unix Shell	4.2%
		T1059.006: Python	3.4%
		T1059.007: JavaScript	1.5%
		T1059.009: Cloud API ☁	0.6%
		T1059.010: AutoHotKey & AutoIT	0.6%
		T1059.005: Visual Basic	0.2%
		T1059.002: AppleScript	0.2%
		T1059.011: Lua ☁	0.2%
T1569: System Services	17.8%	T1569.002: Service Execution	17.8%
T1053: Scheduled Task/Job ☁	13.5%	T1053.005: Scheduled Task	12.7%
		T1053.003: Cron	0.8%
T1204: User Execution ☁	10.4%	T1204.002: Malicious File	6.8%
		T1204.001: Malicious Link	3.6%
T1047: Windows Management Instrumentation	6.3%		
T1559: Inter-Process Communication	0.8%		
T1203: Exploitation for Client Execution	0.4%		
T1072: Software Deployment Tools	0.4%		
T1129: Shared Modules	0.2%		

Bibliography

1. https://www.mandiant.com/resources/blog/attacker-visibility-threat-campaigns
2. https://cloud.google.com/blog/topics/threat-intelligence/analyzing-dark-crystal-rat-backdoor/
3. https://storage.googleapis.com/gweb-uniblog-publish-prod/documents/Year_in_Review_of_ZeroDays.pdf
4. https://cloud.google.com/blog/topics/threat-intelligence/gru-disruptive-playbook
5. https://cloud.google.com/blog/topics/threat-intelligence/chinese-espionage-tactics
6. https://forums.ivanti.com/s/article/CVE-2023-46805-Authentication-Bypass-CVE-2024-21887-Command-Injection-for-Ivanti-Connect-Secure-and-Ivanti-Policy-Secure-Gateways
7. https://cloud.google.com/blog/topics/threat-intelligence/suspected-apt-targets-ivanti-zero-day
8. https://www.ivanti.com/blog/security-update-for-ivanti-connect-secure-and-policy-secure
9. https://services.google.com/fh/files/misc/ivanti-connect-secure-remediation-hardening.pdf
10. https://cloud.google.com/blog/topics/threat-intelligence/ivanti-post-exploitation-lateral-movement
11. https://cloud.google.com/blog/topics/threat-intelligence/mitigating-dprk-it-worker-threat
12. https://www.europol.europa.eu/media-press/newsroom/news/europol-coordinates-global-action-against-criminal-abuse-of-cobalt-strike
13. https://www.cobaltstrike.com/blog/update-stopping-cybercriminals-from-abusing-cobalt-strike
14. https://www.mandiant.com/resources/blog/how-mandiant-tracks-uncategorized-threat-actors
15. https://cloud.google.com/blog/topics/threat-intelligence/apt44-unearthing-sandworm
16. https://cloud.google.com/blog/topics/threat-intelligence/apt45-north-korea-digital-military-machine
17. https://cloud.google.com/security/resources/insights/targeted-attack-lifecycle
18. https://cloud.google.com/blog/topics/threat-intelligence/unc2165-shifts-to-evade-sanctions
19. https://cloud.google.com/blog/topics/threat-intelligence/unc5537-snowflake-data-theft-extortion
20. https://cloud.google.com/blog/topics/threat-intelligence/unc3944-targets-saas-applications
21. https://cloud.google.com/blog/topics/threat-intelligence/unc4393-goes-gently-into-silentnight
22. https://cloud.google.com/blog/topics/threat-intelligence/russia-targeting-signal-messenger
23. https://cloud.google.com/blog/topics/threat-intelligence/unc5537-snowflake-data-theft-extortion
24. https://medium.com/@RadiantCapital/radiant-capital-incident-update-e56d8c23829e
25. https://cloud.google.com/blog/topics/threat-intelligence/mitigating-dprk-it-worker-threat
26. https://www.justice.gov/opa/media/1320156/dl?inline
27. https://ofac.treasury.gov/media/923126/download?inline
28. https://www.justice.gov/usao-dc/media/1352191/dl
29. https://ofac.treasury.gov/media/923126/download?
30. https://cloud.google.com/blog/topics/threat-intelligence/likely-iranian-threat-actor-conducts-politically-motivated-disruptive-activity-against
31. https://research.checkpoint.com/2024/bad-karma-no-justice-void-manticore-destructive-activities-in-israel/
32. https://cloud.google.com/blog/topics/threat-intelligence/unc1860-iran-middle-eastern-networks

33. https://research.checkpoint.com/2024/iranian-malware-attacks-iraqi-government/
34. https://cloud.google.com/blog/topics/threat-intelligence/telegram-malware-iranian-espionage
35. https://www.cybercom.mil/Media/News/Article/2897570/iranian-intel-cyber-suite-of-malware-uses-open-source-tools/
36. https://www.trendmicro.com/en_us/research/24/h/threat-actors-target-middle-east-using-fake-tool.html
37. https://www.gov.il/BlobFolder/reports/alert_1817/he/ALERT-CERT-IL-W--1817.pdf
38. https://cloud.google.com/blog/topics/threat-intelligence/suspected-iranian-unc1549-targets-israel-middle-east
39. https://cloud.google.com/blog/topics/threat-intelligence/untangling-iran-apt42-operations
40. https://blog.google/threat-analysis-group/iranian-backed-group-steps-up-phishing-campaigns-against-israel-us/
41. https://cloud.google.com/blog/topics/threat-intelligence/unc3944-targets-saas-applications
42. https://cloud.google.com/learn/paas-vs-iaas-vs-saas?hl=en
43. https://www.reuters.com/business/finance/goldman-sachs-plans-spin-out-its-digital-assets-platform-bloomberg-news-reports-2024-11-18/
44. https://www.jpmorgan.com/kinexys/digital-payments
45. https://www.blackrock.com/us/financial-professionals/investments/products/bitcoin-investing
46. https://cloud.google.com/blog/topics/threat-intelligence/3cx-software-supply-chain-compromise
47. https://cloud.google.com/blog/topics/threat-intelligence/examining-web3-heists
48. https://www.justice.gov/archives/opa/pr/911-s5-botnet-dismantled-and-its-administrator-arrested-coordinated-international-operation
49. https://www.virustotal.com/gui/collection/campaign--a35afe10-cd7e-5c2c-b2d4-21cdaf3d9a75
50. https://cloud.google.com/blog/topics/threat-intelligence/securing-cryptocurrency-organizations

Google Cloud

For more information, visit **cloud.google.com**.