Mandiant

Google Cloud Security

M-Trends

2026 Report

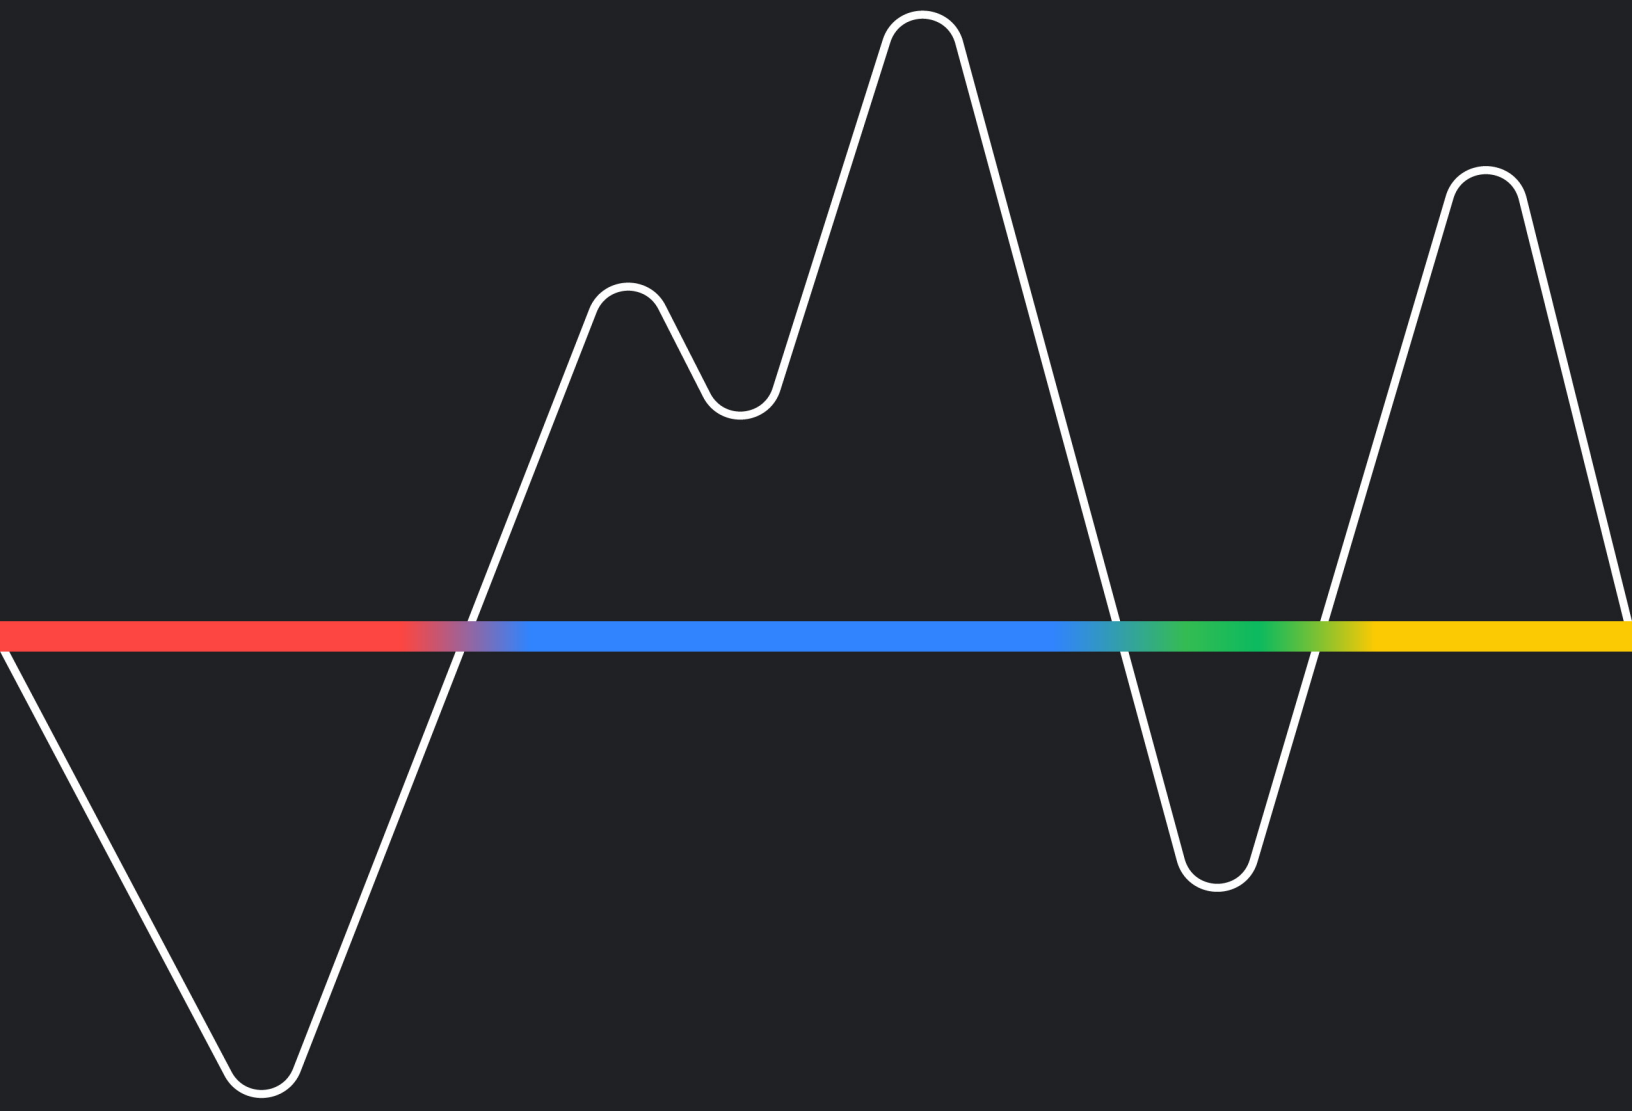

Table of Contents

Introduction

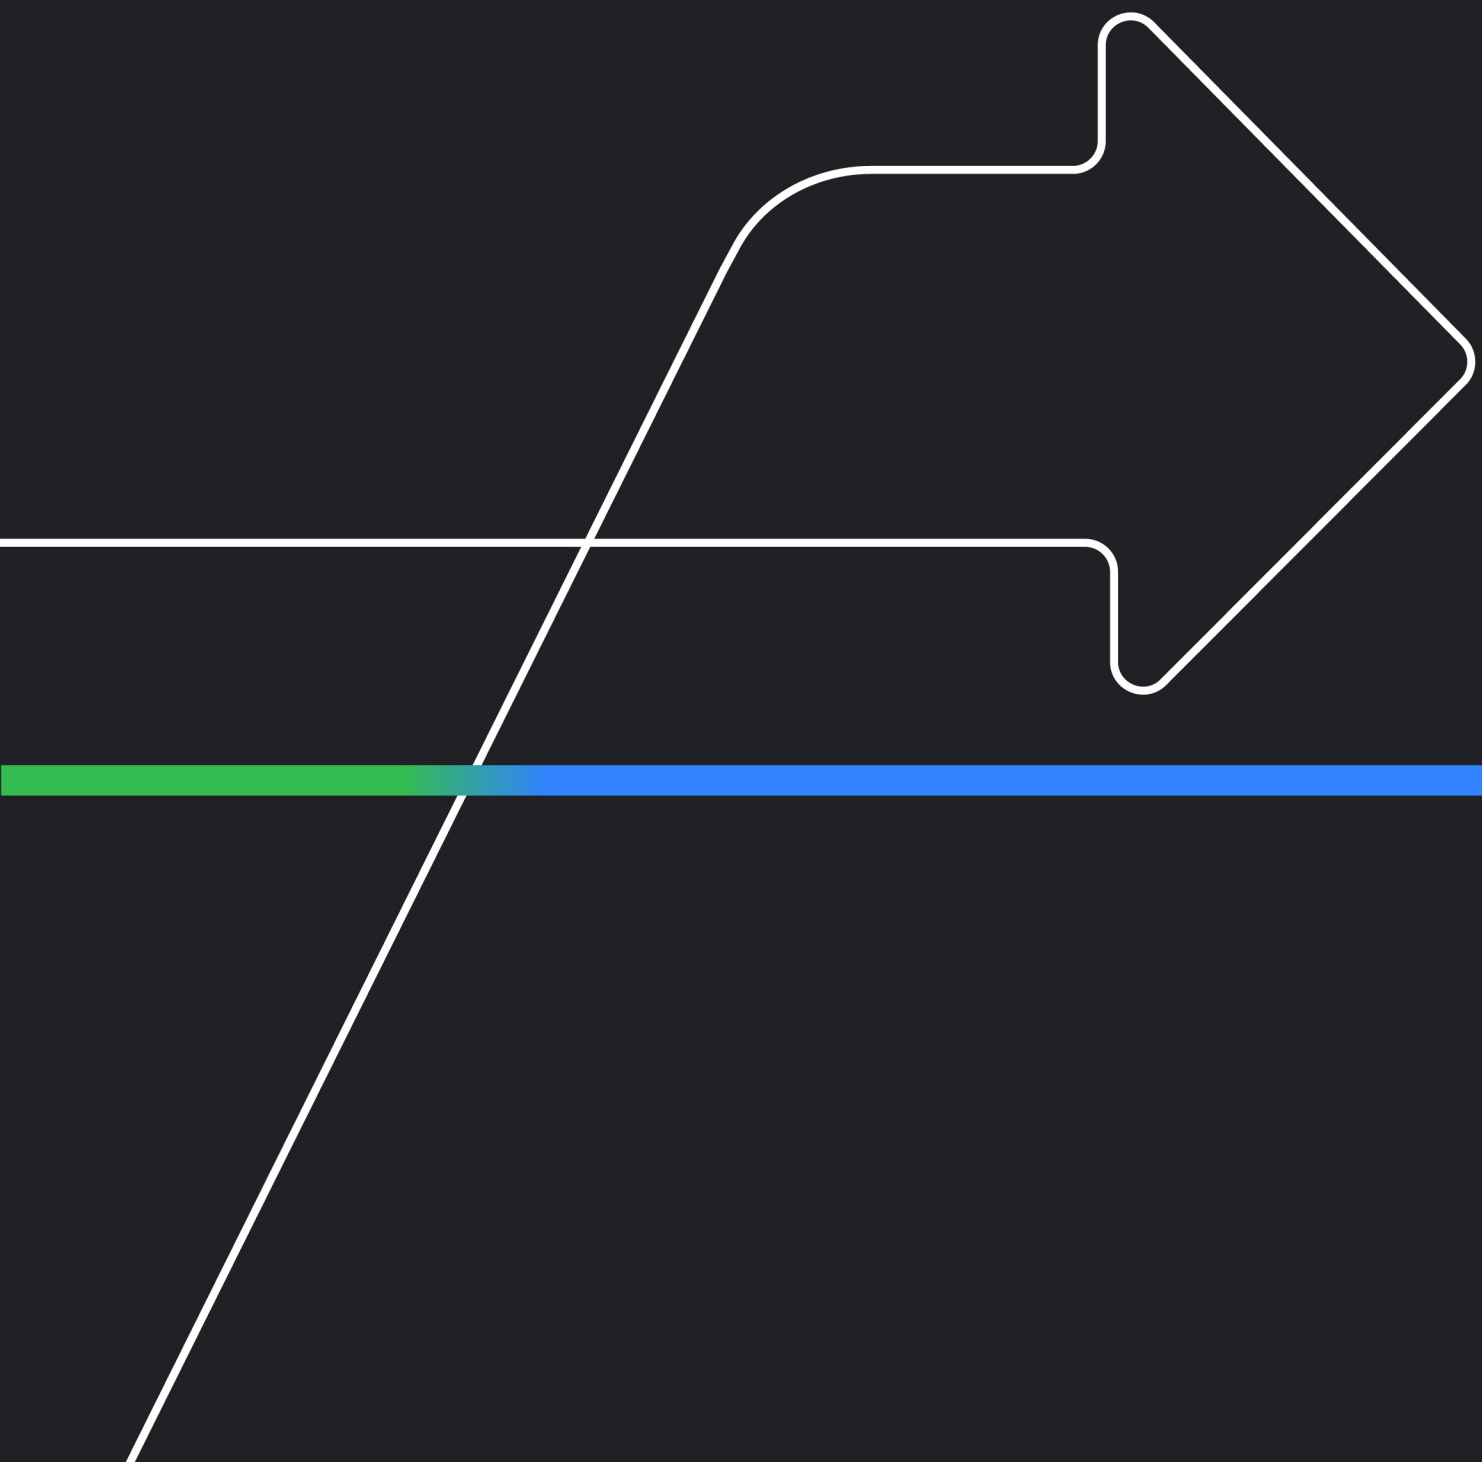

M-Trends serves as a definitive look at the threats and tactics used in breaches, grounded in over 500k hours of frontline incident investigations conducted by Mandiant in 2025. Together with Google Threat Intelligence Group (GTIG), we have a comprehensive view of the modern threat landscape, and emerging threats that may drive future attacks.

One of the dynamics influencing the threat landscape is AI. Recent GTIG reporting confirms that state-sponsored and financially motivated actors are integrating AI to accelerate the attack lifecycle. Attackers are increasingly relying on large language models (LLMs) as a strategic force multiplier to move beyond mass email campaigns toward hyper-personalized, rapport-building social engineering. In the wild, malware families like PROMPTFLUX and PROMPTSTEAL actively query LLMs mid-execution to evade detection, while "distillation attacks" threaten intellectual property by extracting the proprietary logic and specialized training data of high-value machine learning models.

To ensure organizations are prepared for these evolving capabilities, Mandiant red teams are actively incorporating these AI-driven techniques into engagements, rigorously testing defenses against emerging threats. However, despite threat actors increasingly leveraging AI, especially during the early phases of the attack lifecycle, we don't consider 2025 to be the year where breaches were the direct result of AI.

From our view on the frontlines, the vast majority of successful intrusions still stem from human and systemic failures. M-Trends 2026 contains these critical insights that organizations must think about today, alongside the actionable guidance required to stay ahead of modern threats.

We are tracking a significant shift toward voice-based social engineering (vishing), which has risen to the number two spot for initial infection vectors. Exploits remain the most common vector for a sixth consecutive year. Furthermore, our investigations highlight a rising "hand-off" trend, where one threat actor gains initial access, and then provides access to a separate actor that typically conducts higher impact operations like ransomware. This hand-off is happening so fast (sometimes under 30 seconds) that it creates a scenario where alerts traditionally considered "lower priority" can very quickly become significant compromises.

Other M-Trends 2026 observations include:

- Global median dwell time has risen to 14 days from 11 days, driven largely by long-term espionage and DPRK IT worker operations. The longer adversaries remain undetected, the greater the cost and complexity of remediation.
- Financially motivated and cyber espionage groups continue to abuse native functionalities in on-premises and cloud environments, as well as legitimate tools, to reduce opportunities for detection. This means traditional endpoint security relying on malware signatures is not sufficient.
- Ransomware operators have shifted their primary objective from data theft to deliberate recovery denial, systematically targeting backup infrastructure, identity services, and virtualization management planes. By destroying the ability to recover, threat actors put immense pressure on organizations to pay ransom demands.

Mandiant converts frontline observations into the strategic intelligence required to help organizations fortify their defenses. To defend against the activity in M-Trends 2026, organizations should prioritize the security of infrastructure such as backups, identity services, and the virtualization layer, which attackers are now systematically targeting to deny recovery. Hardening edge and core network devices remains critical, especially since exploits continue to be the most common entry point for adversaries. By addressing these specific visibility gaps and focusing on the tactics that actually bring attackers success, defenders will strengthen their cyber resilience.

Mandiant consultants operate globally on the frontlines of the latest cyber incidents, observing the most relevant adversary tactics, techniques, and procedures (TTPs). This perspective provides us with a unique understanding of the modern threat landscape, and the strategies required to defend against it. We translate this intelligence into action by proactively hardening client environments against these TTPs, and helping them feel confident in their cybersecurity readiness. The information in this report has been sanitized to protect the identities of targeted organizations and their data.

By the Numbers

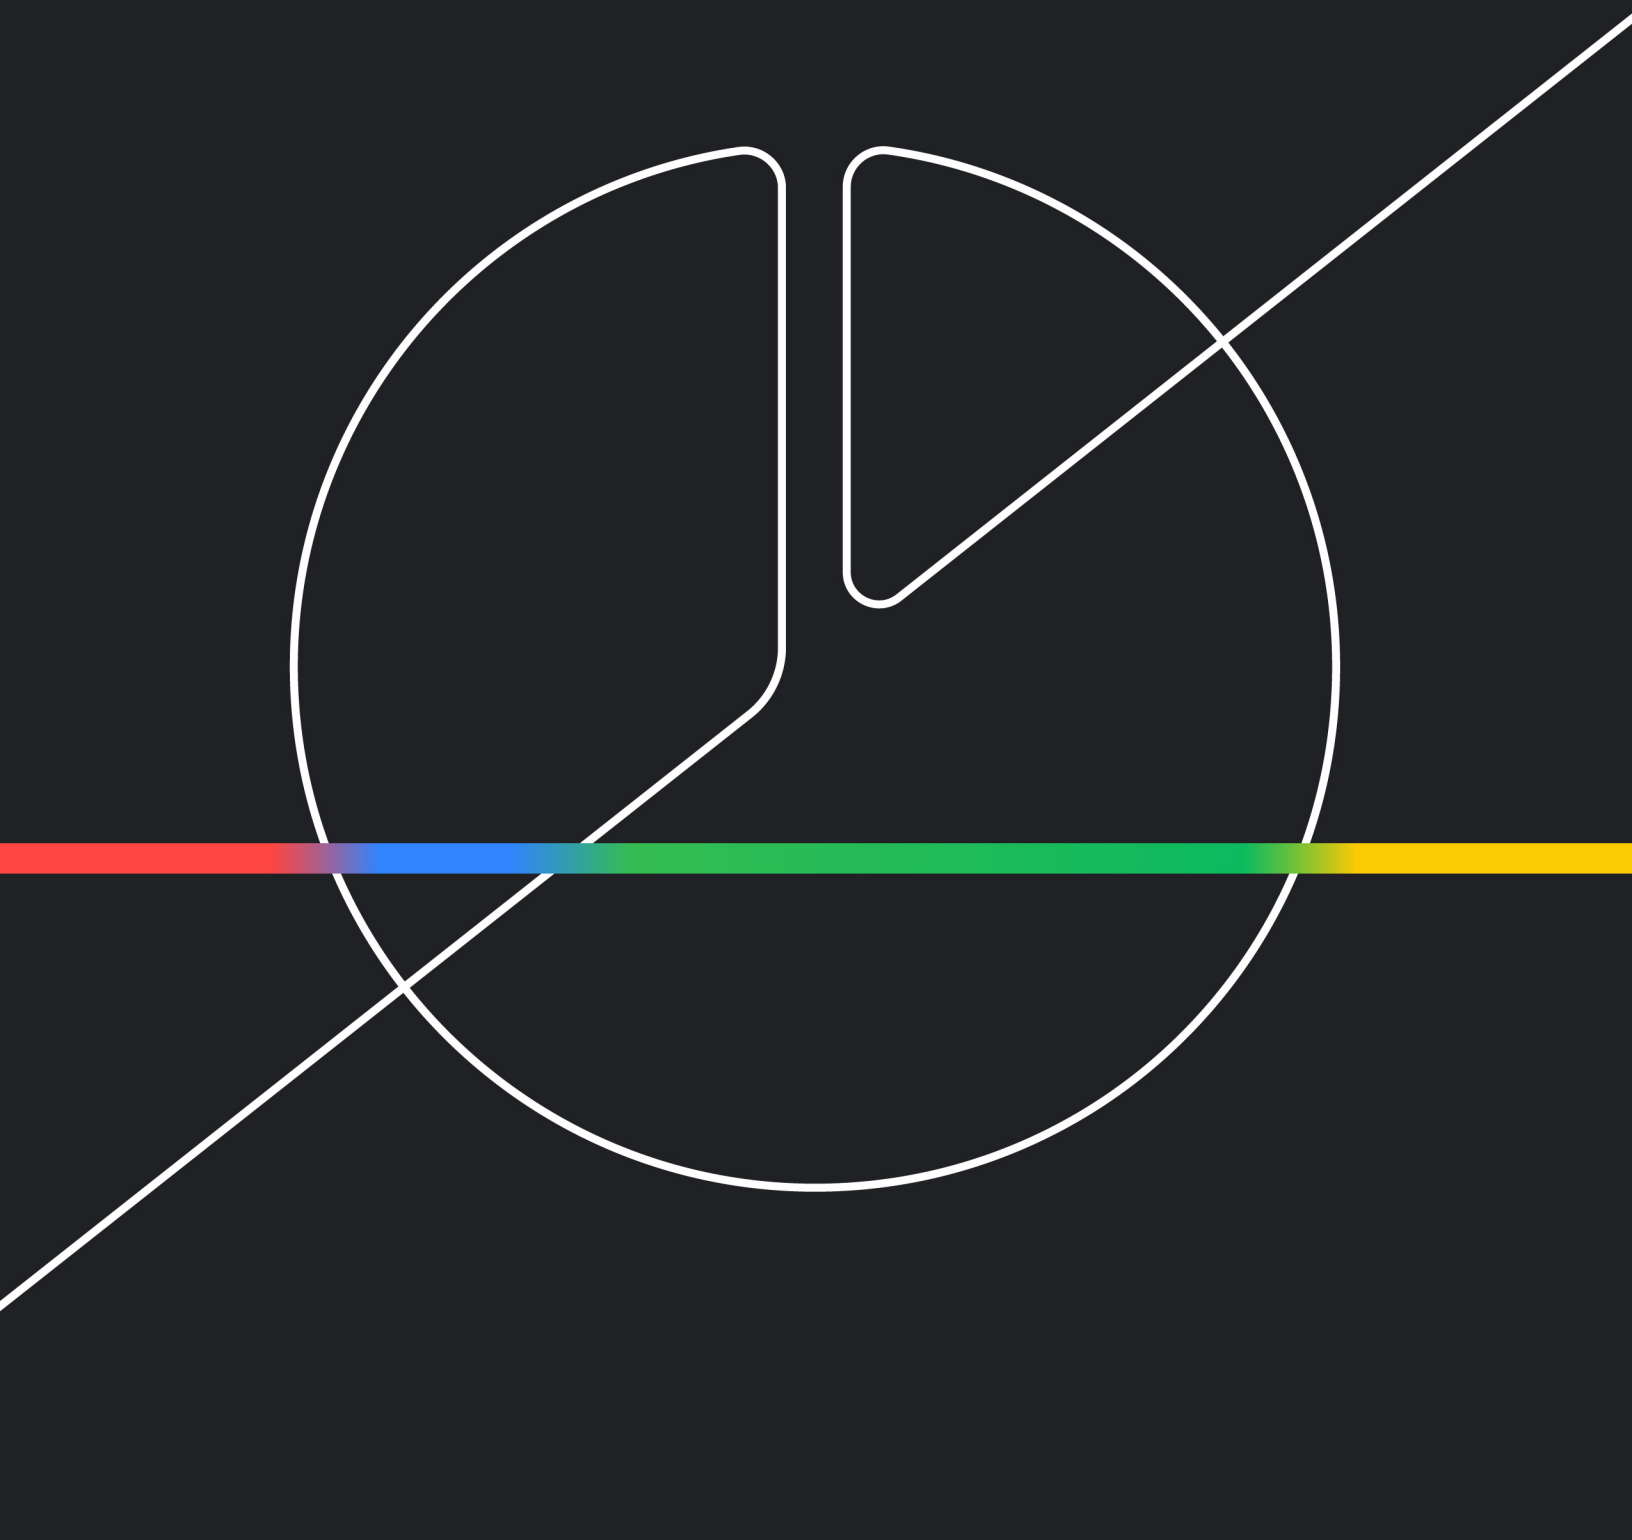

Since 2010, Mandiant has provided statistics and analysis of threats observed in the previous year's incident investigations. In M-Trends 2026, Mandiant examines data collected from more than 500,000 hours of incident response engagements globally, highlighting trends and significant insights.

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 contained in By The Numbers are based on targeted attack activity identified by Mandiant between Jan. 1 and Dec. 31, 2025.

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.

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

To provide a comprehensive view of the global threat landscape, Google Threat Intelligence Group (GTIG) tracks impactful activity through two distinct categories: Campaigns and Global Events. Campaigns represent focused efforts by one or more threat groups pursuing a single objective, while Global Events encompass multiple groups using similar tactics, such as exploiting a new vulnerability. These records are continuously updated with new indicators of compromise (IOCs), tactics, techniques, and procedures (TTPs), and affected industries and regions. Whenever possible, GTIG provides examples, context, and information about threat cluster 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 GTIG research, enabling our clients to respond effectively and decisively to active threats at first discovery and as they evolve.

In 2025, GTIG initiated 83 campaigns and eight global events and continued to track activity identified in previous years. These campaigns affected every industry vertical and 73 countries across six continents based on GTIG's direct observations. A subset of this activity—specifically 35 campaigns and six global events—include incidents that Mandiant investigated. Insights drawn from the investigations provide deep forensic insight into these events and broader trends.

The diagram below summarizes the Campaigns and Global Events. Prominent examples include the exploitation of CVE-2025-31324 in SAP NetWeaver, and the data theft extortion campaign exploiting a zero-day vulnerability and using GOLDVEIN.JAVA.

One of the most active global events was the growing adoption of the ClickFix social engineering technique, which involves clusters using a prompt, typically on a phishing page, to convince users to execute PowerShell or other system-level commands under the guise of fixing problems or verifying their legitimacy. Variants of the technique observed in 2025 used a variety of lures, including CAPTCHAs, verifications for video conference and meeting invitations, driver or operating system updates, and enterprise software compliance verification. In 2025, GTIG identified dozens of threat clusters incorporating this technique, particularly threat clusters focused on widespread initial access operations.

2025 Campaigns and Global Events Related to Mandiant Incident 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.

In 2025, Mandiant responded to incidents in the high-tech sector most frequently, followed by the financial sector, business and professional services, and healthcare. These sectors consistently appear at the top of the list for industry distribution, though in 2025 investigations in the high-tech sector outpaced those of financial sector organizations, which had the largest share of Mandiant investigations in 2024 and 2023.

Targeted Industries, 2025

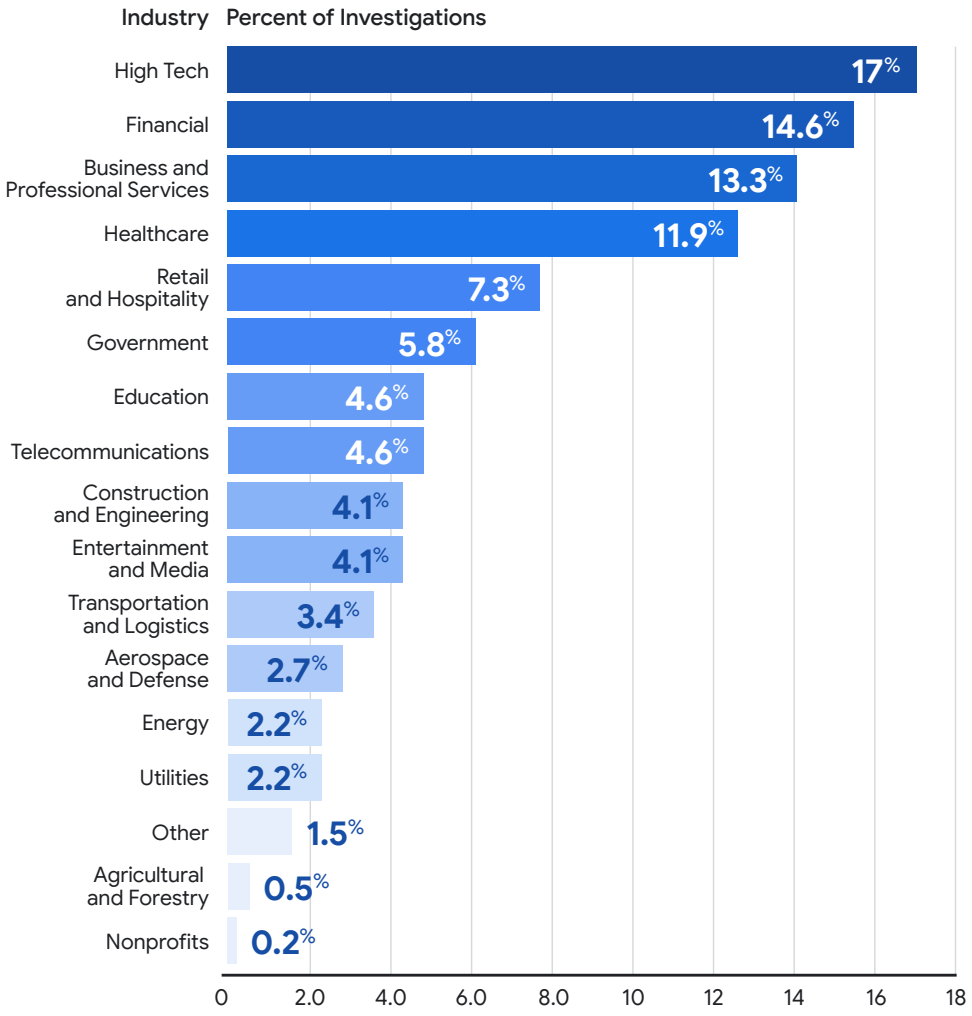

Initial Infection Vector

For the sixth year running, exploits represented the most frequently observed initial infection vector in 2025 Mandiant incident response investigations. Exploits comprised 32% of investigations in which an initial infection vector could be identified.

Initial Infection Vector, 2025

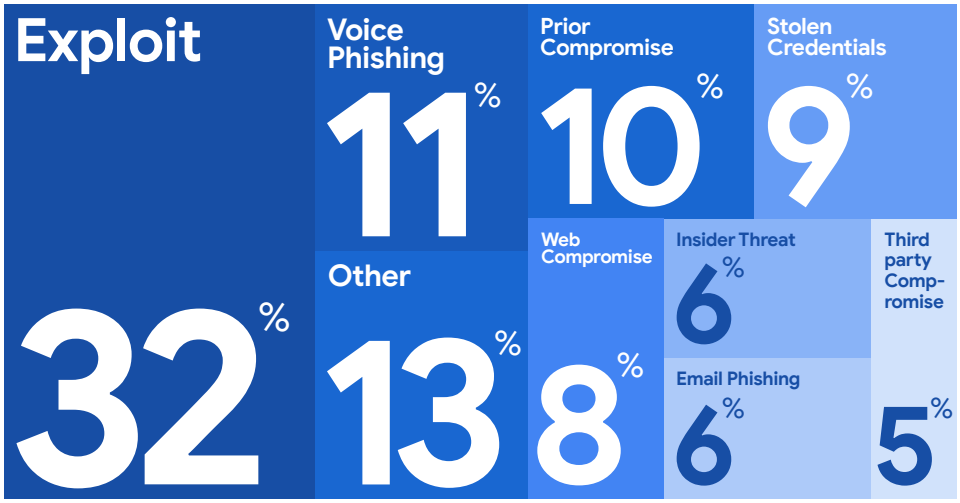

In 2025, threat clusters used an increasingly diverse array of social engineering tactics across email, voice, messaging platforms, and social media. To capture this nuance, GTIG has refined these categories to distinguish between interactive human engagement, such as voice phishing and non-interactive technical lures, such as email phishing. While email phishing often relies on volume and opportunistic delivery, interactive methods involve a live person steering the conversation in real-time. This distinction is critical for defenders: interactive attacks are significantly more resilient against automated technical controls and require different detection strategies.

Globally, email phishing was no longer a top observed initial intrusion vector. At 6% of 2025 intrusions, it continues to see a steady decline from its 14% share in 2024. Conversely, the proportion of incidents stemming from social engineering through voice phishing and messaging apps noticeably increased in 2025 compared to prior years, with voice phishing emerging as the second-most commonly observed vector in 2025 investigations at 11%.

One of the more pervasive examples of this activity was a campaign that spanned the first half of 2025, in which UNC6040 used voice phishing to convince targets to provide credentials and authorize an attacker-controlled version of a legitimate software-as-a-service (SaaS) application to access organizations' data. These organizations later received ShinyHunters-branded extortion notes demanding payment for the non-release of stolen data. Given the significant time lapse between the initial data theft activity and the extortion operations, GTIG tracks the extortion activity as UNC6240. Another example of a long-term voice phishing campaign came from UNC3944, a financially motivated threat cluster that has been active since at least early 2022 and overlaps with public reporting on Scattered Spider. UNC3944 targeted help desk staff by impersonating employees requesting password resets and changes to multi-factor authentication (MFA) settings.

Email Phishing Declines as an Initial Infection Vector, 2022-2025

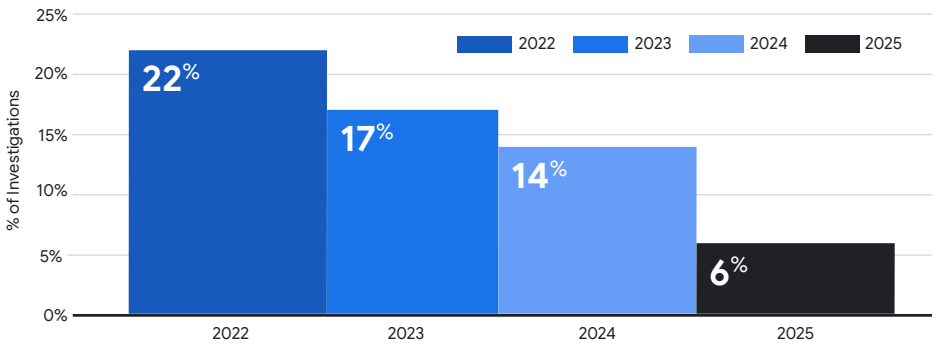

Though there was a decline in email phishing as an initial infection vector, email phishing continues to be a consistently leveraged tactic. For example, UNC6345 used a phishing email with malicious employee rewards and benefits-themed PDF to compromise a user. Successful email phishing has also been known to extend the reach of a compromise; for example, UNC3203 gained access to an Microsoft 365 environment, created mailbox forwarding rules to forward received emails to an external address, then leveraged the access to send phishing emails internally and externally. In a similar case, another threat cluster used phishing emails containing malicious attachments disguised as financial documents to deploy malware that uses Outlook COM automation to hijack targeted users' mailboxes and facilitate lateral movement.

Prior compromise was the third-most common initial infection vector across investigations performed in 2025. In 2024, prior compromise was the fifth-most common form of initial infection, found in 8% of investigations. Mandiant identified stolen credentials as the initial infection vector in 9% of 2025 investigations, which is a decline from 16% in 2024. Web compromise remained relatively stable at 8% of investigations.

The share of incidents traced back to malicious insiders increased from 5% in 2024 to 6% in 2025. Consistent with 2024, the majority of these insiders were North Korean IT workers, though Mandiant also identified evidence of financially motivated threat clusters bribing employees or contractors for their corporate credentials.

Third-party compromise represented 5% of identified initial infection vectors in 2025 investigations. This figure includes a number of incidents involving compromises of software-as-a-service (SaaS) platforms. Additional initial infection vectors included software supply chain compromise, brute force, server compromise, ClickFix, and SQL injection.

For **guidance regarding protecting credentials** and securing against common initial infection vectors including stolen credentials, please see:

Keys to the Kingdom: A Defender's Guide to Privileged Account Monitoring

For additional resources about **voice phishing threats**, please see:

Hello, Operator? A Technical Analysis of Vishing Threats

The Cost of a Call: From Voice Phishing to Data Extortion

Initial Infection Vectors

Exploits include initial access achieved by exploitation of a software vulnerability with a CVE identifier.

In **social engineering** operations, threat clusters attempt to trick targets into supplying credentials, downloading malware, or otherwise providing the attacker with access to sensitive information and/or systems. Commonly referred to as **phishing**, this can take place over **email**, the **phone**, **chat** or messaging applications, **SMS**, or **social media**.

Prior compromise refers to incidents in which a threat cluster gained access to a targeted environment through a foothold established by a separate threat cluster.

Stolen credentials can be sourced directly from infostealers, harvested from inadvertently exposed databases or source code repositories, or indirectly gathered from dark web forums or database leaks. These typically include incidents in which the first evidence of malicious activity is the threat actor logging on using valid credentials.

Web compromise encompasses drive-by compromise, the use of malicious advertisements, search engine optimization (SEO) poisoning, and compromised websites.

Insider threat describes compromises traced back to people within the organization, such as employees, former employees, contractors, or business associates who abuse or misuse their access, leading to attackers gaining information about an organization's security practices, data, and computer systems.

Third-party compromises occur when an attacker gains unauthorized access to accounts or infrastructure belonging to one organization and then uses that foothold to gain access to additional targets.

Most Frequently Exploited Vulnerabilities

The most frequently exploited vulnerabilities identified in 2025 Mandiant incident response investigations were zero-days affecting internet-facing web application servers. These vulnerabilities, either alone or chained with additional flaws, enabled unauthenticated code execution against enterprise platforms that provide centralized access to an organization's financial data, business operations data, or internal documents. Threat clusters often see these types of targets as opportunities for reconnaissance and a beachhead from which they can expand further into a compromised network.

Most Frequently Exploited Vulnerabilities

CVE-2025-31324

CVE-2025-31324 is an improper authorization vulnerability in SAP NetWeaver Visual Composer, specifically in its Metadata Uploader component, allowing unauthenticated attackers to upload arbitrary files. This flaw can lead to unauthorized access and be chained with another vulnerability, such as CVE-2025-42999, to achieve code execution.

Mandiant responded to a number of incidents in which threat actors exploited CVE-2025-31324. GTIG observed evidence that at least four separately tracked threat clusters likely exploited CVE-2025-31324 as a zero day in early 2025. After SAP issued a patch in April 2025, GTIG tracked six additional threat clusters, including several suspected PRC-nexus cyber espionage clusters, exploiting the vulnerability as an n-day. The activities observed post compromise were limited to attackers establishing a foothold in targeted environments with web shells or backdoors and conducting reconnaissance.

For more information about the **CVE-2025-61882 exploitation**, please see:

Oracle E-Business Suite Zero-Day Exploited in Widespread Extortion Campaign

CVE-2025-61882

CVE-2025-61882 is an improper authentication vulnerability affecting Oracle E-Business Suite (EBS) that allows an unauthenticated, remote attacker with network access to achieve arbitrary code execution. Mandiant investigated multiple incidents involving the exploitation of CVE-2025-61882.

In September 2025, a threat cluster claiming affiliation with the CL0P extortion brand sent extortion emails claiming that they had compromised organizations Oracle E-Business Suite (EBS) applications and stolen documents. In related incident response engagements, Mandiant identified evidence of attempted exploitation of Oracle EBS dating back to July 2025, and successful exploitation in August 2025 that may constitute zero-day exploitation of CVE-2025-61882. Following exploitation, the threat actors deployed Java payloads, including GOLDVEIN.JAVA.

GTIG attributes this activity to a suspected FIN11 threat cluster based on several similarities with past FIN11 campaigns, including the use of the CL0P data leak site (DLS). The in-memory Java-based loader GOLDVEIN.JAVA that fetches a second-stage payload is reminiscent of the GOLDVEIN downloader and GOLDTOMB backdoor observed during the mass exploitation of the Cleo Managed File Transfer (MFT) vulnerability in late 2024, which was attributed to another suspected FIN11 cluster. More broadly, exploitation of a zero-day vulnerability in a widely used enterprise application, followed by a large-scale, branded extortion campaign weeks later is a hallmark of activity historically attributed to FIN11.

CVE-2025-53770

Microsoft SharePoint Server was affected by a deserialization of untrusted data vulnerability, CVE-2025-53770, which allows an unauthenticated, remote attacker to execute arbitrary code on the affected SharePoint server. This flaw can be chained with CVE-2025-53771 in an exploit known as ToolShell. CVE-2025-53771 is a path traversal and spoofing vulnerability that enables unauthorized access to SharePoint content, internal configurations, and system files, as well as the deployment of web shells and other post-exploitation activities.

Mandiant investigated several incidents involving exploitation of these vulnerabilities, and GTIG tracked evidence of widespread exploitation, including at least two threat clusters exploiting CVE-2025-53770 as a zero day, with an additional three threat clusters observed exploiting the flaws after the July 20 and 21 patches were released. Most observed post-compromise activity, including by suspected PRC-nexus cyber espionage threat cluster UNC6349, was focused on reconnaissance and establishing a foothold. In contrast, the financially motivated threat cluster UNC6357 exploited the SharePoint vulnerabilities to ultimately deploy the LOCKBIT.WARLOCK ransomware.

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 entity notification is when an outside entity informs an organization it has been compromised, such as law enforcement agencies, cybersecurity companies, or industry partners.

In 2025, 52% of organizations detected evidence of malicious activity internally. External entities, such as law enforcement, CERTs, or cybersecurity companies, notified organizations of a potential compromise in 34% of cases. Adversaries informed organizations of a compromise, typically in the form of a ransom note, in 14% of cases. The proportion of internally detected compromises increased from 43% in 2024 to 52% in 2025, while external entity notifications declined from 43% to 34%.

Global Detection by Source, 2025

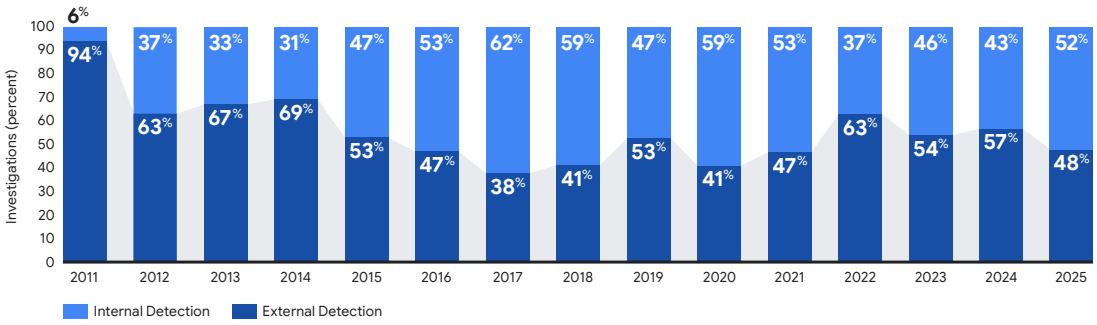

Global Detection by Source, 2025

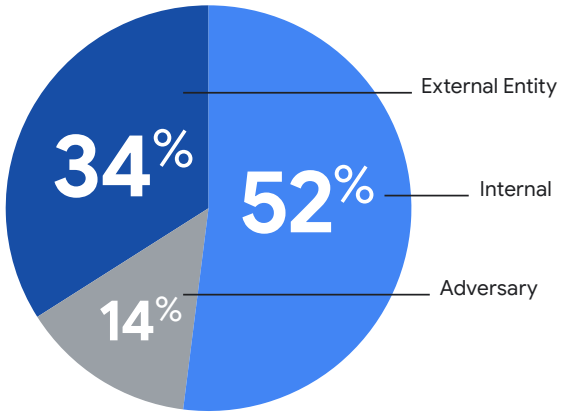

Global Median Dwell Time

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.

Global median dwell time across 2025 Mandiant investigations was 14 days. This is an increase from 11 days in 2024. Median dwell time for incidents discovered internally in 2025 remained fairly consistent with 2024 at 9 days. Organizations were made aware of an incident by an external notification in 25 days in 2025, a significant increase in the external notification global median dwell time from 2024, 11 days.

Median Dwell Time, 2011-2025

	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025
All	416	243	229	205	146	99	101	78	56	24	21	16	10	11	14
External	—	—	—	—	320	107	186	184	141	73	28	19	13	11	25
Internal	—	—	—	—	56	80	57.5	50.5	30	12	18	13	9	10	9

Median Dwell Time by Detection Source, 2025

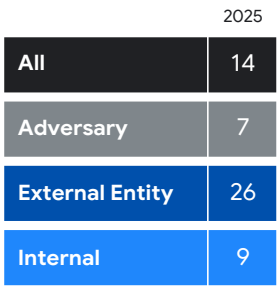

Median dwell time for incidents in which adversaries notified organizations of a compromise, often in the form of a ransom demand, was seven days in 2025, slightly higher than five days in 2024.

Global Dwell Time Distribution

Comparing dwell time distributions from 2024 to 2025 reveals a slight shift: very short dwell times decreased, while intermediate dwell times (one week to six months) saw a marginal increase. There was a small decrease in the share of incidents discovered in a week or less, down from 45.1% to 41.5%. The share of incidents that went undiscovered for more than a week to a month, showed a small increase from 17.6% to 20.1%. A similar increase was observed for incidents that remained undetected for one to six months, growing from 23.9% to 26.7%.

This observed shift toward longer dwell times likely reflects the quantity of incidents in which threat clusters prioritize maintaining long-term access to targeted environments, including cyber espionage, North Korean IT workers, and other types of compromises. Mandiant has directly observed these types of groups make concerted efforts to remain undetected. They leverage living-off-the-land (LotL) techniques, minimize use of custom malware, remove artifacts, favor obfuscation, and mimic legitimate products or system tools already present within the victim's environment. The median dwell times for incidents assessed to be motivated by cyber espionage, as well as North Korean IT worker incidents, were both 122 days, or about four months.

The percentages of incidents discovered more than six months after the first evidence of malicious activity declined. Despite the shifts in the distribution of dwell times across our datasets from 2024 to 2025 revealing more dwell times in the intermediate range in 2025, the multi-year comparison of dwell time distribution continues to indicate that, in the long term, dwell times are getting shorter.

Global Dwell Time Distribution, 2018-2025

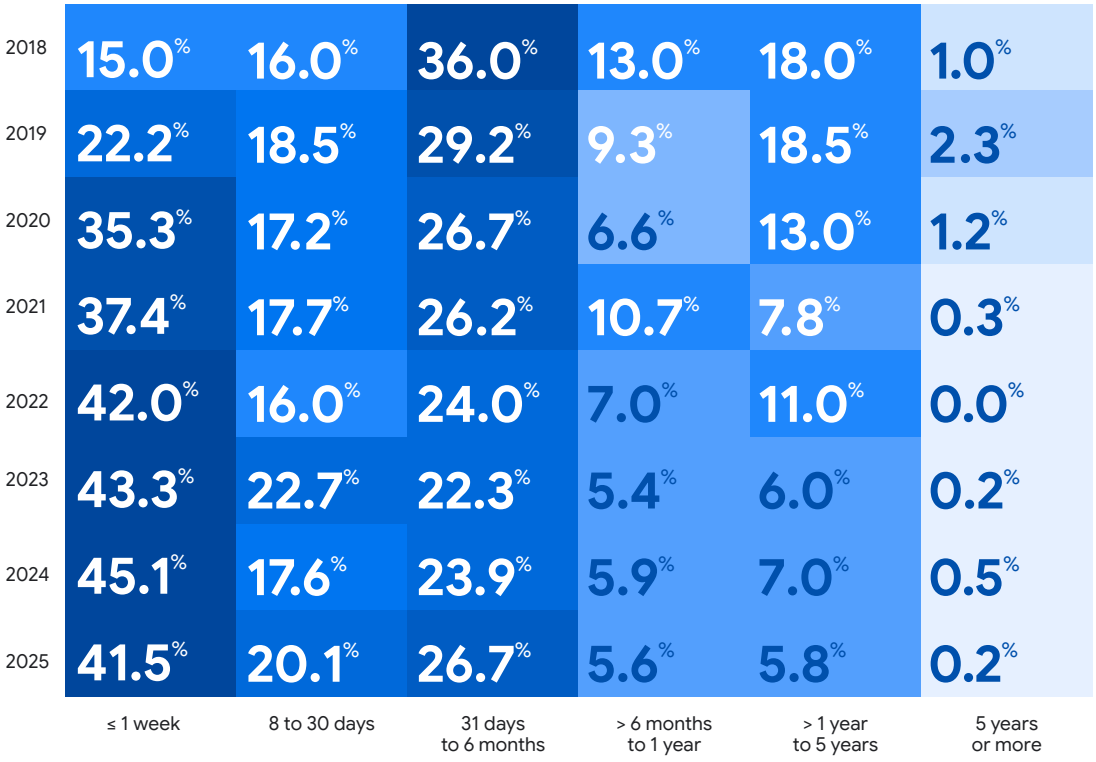

Post-Compromise Activity

Financial Gain, 2020-2025

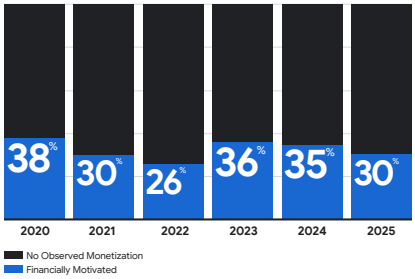

2025 Financial Gain Incidents by Type

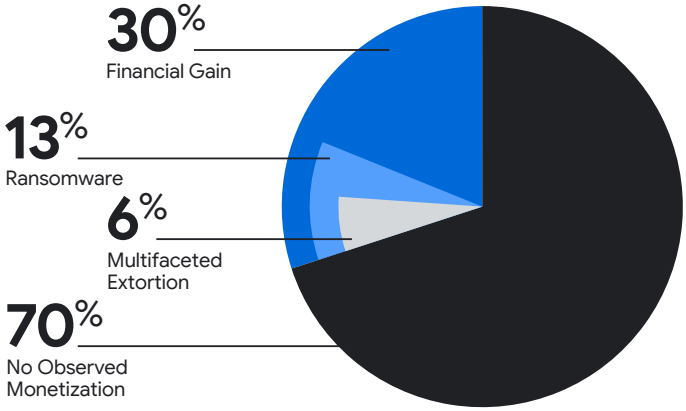

Financial Gain

Threat actors used monetization techniques in 30% of the investigations that Mandiant performed in 2025. This is a decline from 35% of 2024 incidents. Extortion-related intrusions, which includes ransomware as well as data theft extortion without ransomware encryption, represented 23% of 2025 intrusions and approximately three fourths of financially motivated intrusions. Intrusions in which the threat actors attempted to or successfully deployed ransomware comprised 13% of the incidents Mandiant investigated in 2025.Multifaceted extortion, or incidents involving both ransomware encryption and data theft extortion, constituted 6% of 2025 compromises.

In addition to extortion-related incidents, Mandiant encountered threat clusters pursuing a variety of other monetization methods, including North Korean IT worker employment fraud, payment redirection fraud, selling access to compromised networks, ATM malware, use of web skimmers to capture credit card information, as well as theft of cryptocurrency, loyalty points, and gift card data.

Data Theft, 2020-2025

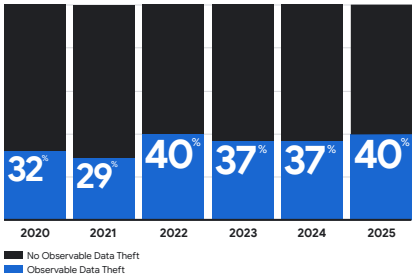

2025 Data Theft Incidents by Type

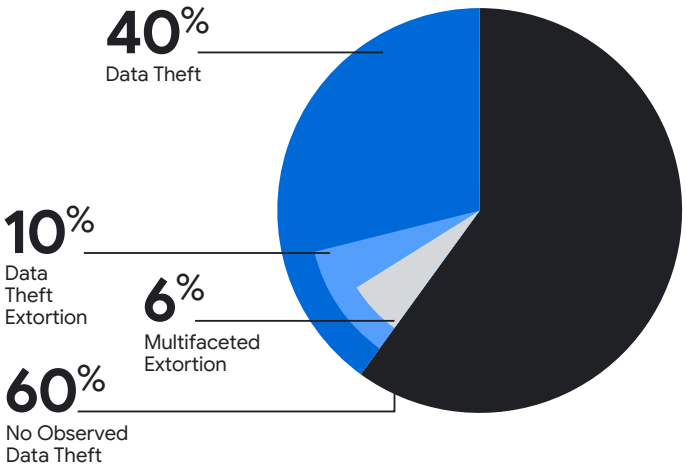

Data Theft

Mandiant identified evidence of data theft in 40% of investigations performed in 2025. This is slightly higher than the proportion from 2024, 37%. Incidents that involved data theft extortion constituted 10% of 2025 investigations, which is on par with 2024, when data theft extortion events represented 10% of all investigations.

In many of the 2025 investigations in which Mandiant identified evidence of data theft, threat actors targeted credentials and reconnaissance data useful for maintaining persistence, lateral movement, and escalating privileges. Other examples of data theft appeared to be wholesale and opportunistic. In several cases, threat actors targeted personally identifiable information (PII), such as customer records listing contact and order information. Mandiant identified threat clusters that used stolen PII in subsequent voice phishing attempts.

For more information about UNC6395, please see:

Widespread Data Theft Targets Salesforce Instances via Salesloft Drift

Mandiant responded to several incidents in which the threat actors targeted code repositories that they subsequently mined for credentials and keys, including several compromises attributed to UNC6395. Mandiant identified several cyber espionage groups, including a suspected UNC5221 cluster; the PRC-nexus cyber espionage cluster UNC5807, which overlaps with the publicly reported "Salt Typhoon"; and a suspected APT44 threat cluster, compromise third-party service providers to facilitate data theft from customers of these service providers. Mandiant also identified evidence of selective data theft. For example, in multiple incidents attributed to disparate PRC-nexus cyber espionage groups, including UNC5221, Mandiant identified evidence that the threat clusters targeted particular users or subjects of interest. While UNC5221 has been used synonymously with the actor publicly reported as Silk Typhoon, GTIG does not currently consider the two clusters to be the same.

UNC5221 Targeting Third-Party SaaS Providers

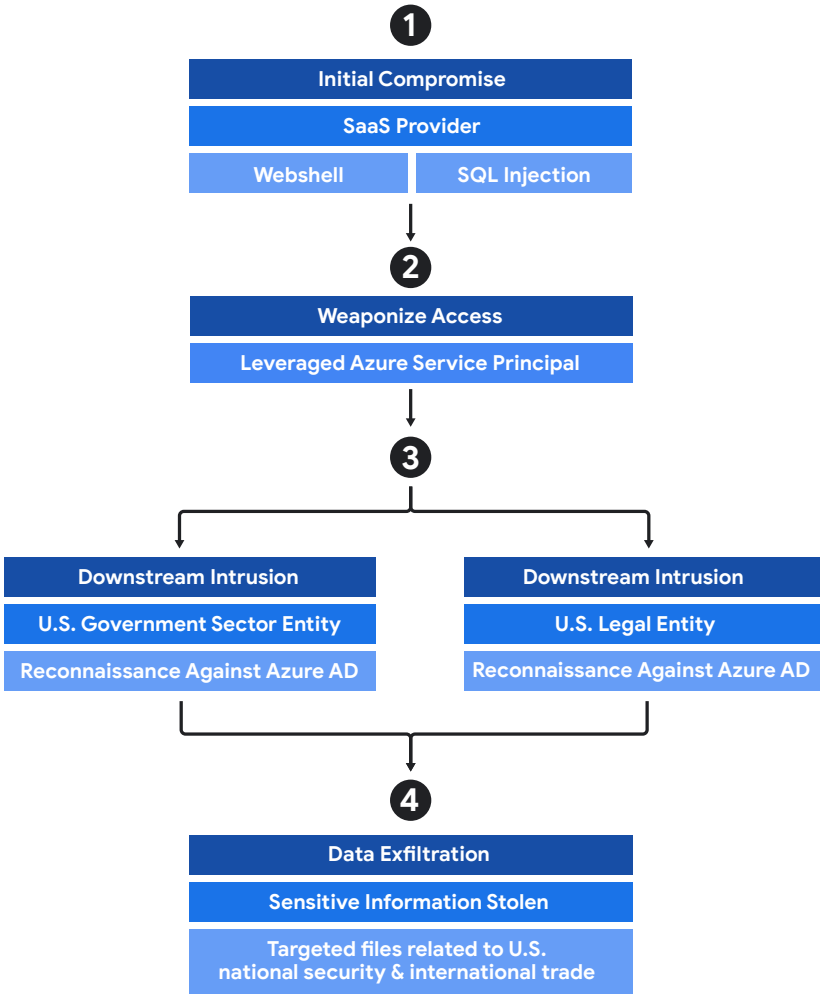

Insider Threat

For more information about North Korean IT Workers, please see:

The Ultimate Insider Threat: North Korean IT Workers

DPRK IT Workers Expanding in Scope and Scale

Staying a Step Ahead: Mitigating the DPRK IT Worker Threat

In 2025, Mandiant responded to a number of incidents involving North Korean IT workers using false or stolen identities to carry out employment fraud and ultimately provide revenue for the North Korean regime. The IT workers Mandiant identified targeted organizations in a variety of industries, including the business and professional services, financial, government, technology, healthcare, and hospitality sectors. Many of the IT workers remained undetected in target environments for long periods of time, in several cases more than one year. The median dwell time for the IT worker incidents Mandiant responded to was 122 days, or about four months.

Mandiant also identified financially motivated threat clusters bribing contractors to provide corporate credentials or other access into targeted organizations, leading to data theft and attempted extortion.

Malware

A malware family is a program or set of associated programs with sufficient "code overlap" among the variants that Mandiant considers 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.

GTIG began tracking 714 new malware families in 2025, a significant increase from 632 in 2024, which brought the total number of tracked malware families to more than 6,000. A total of 224 malware families were observed in the investigations Mandiant performed in 2025, which includes 126 newly tracked families and an additional 98 malware families that were first discovered in prior years. For comparison, GTIG observed 205 malware families and 83 malware families were both newly tracked and observed in 2024 investigations.

As with prior years, the majority of newly tracked malware (72%) and malware families observed (63%) in 2025 investigations were effective on Windows. These percentages are consistent with 2024 findings. Malware families that are effective exclusively on Linux accounted for 12% of newly tracked families and 11% of observed malware families. The percentage of newly tracked malware families effective only on Linux remained stable compared to 2024, but the percent of observed malware families effective only on Linux declined slightly from 17% in 2024. In addition to Windows and Linux, GTIG tracked malware families effective on MacOS, BSD, and Unix.

Tracked and Observed Malware Families in 2025

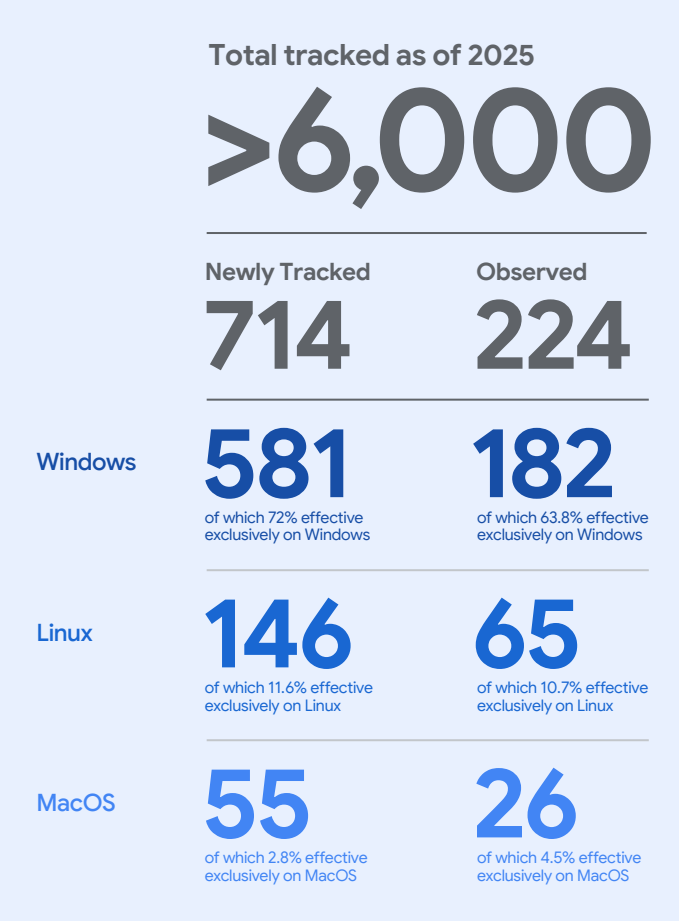

Malware Families by Category

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.

Newly Tracked Malware

Newly tracked malware are families that were first identified and named between Jan. 1 and Dec. 31, 2025. Of these families, 33% were backdoors, 14% were droppers, 14% were downloaders, 6% were ransomware, 6% were launchers, 5% were credential stealers, and 5% were data miners. The "Other" category includes keyloggers, tunnelers, utilities, rootkits, ATM malware, disruptive malware, and point-of-sale malware, each of which make up less than 5% of the malware families observed in 2025. Compared to 2024, there were fewer downloaders (19% in 2024), but other proportions remained consistent with findings from prior years.

Newly Tracked Malware Families by Category, 2025

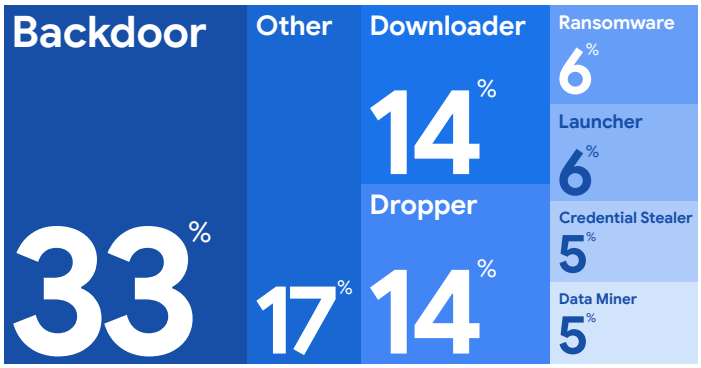

Observed Malware

Observed malware encompasses both legacy and newly tracked malware observed in engagements. This provides a comprehensive look at which malware categories remain the most persistent choices for attackers in real-world environments. The most frequently observed roles of malware families observed in Mandiant's 2025 investigations included backdoors at 36%, followed by downloaders at 11%, ransomware at 10%, droppers at 10%, credential stealers at 9%, with 24% of observed malware families falling into other roles. Despite an increase in newly tracked ransomware families and variants, ransomware declined in its share of observed malware from 14% in 2024, while downloaders increased from 7%, and credential stealers increased from 5%.

Observed Malware Families by Category, 2025

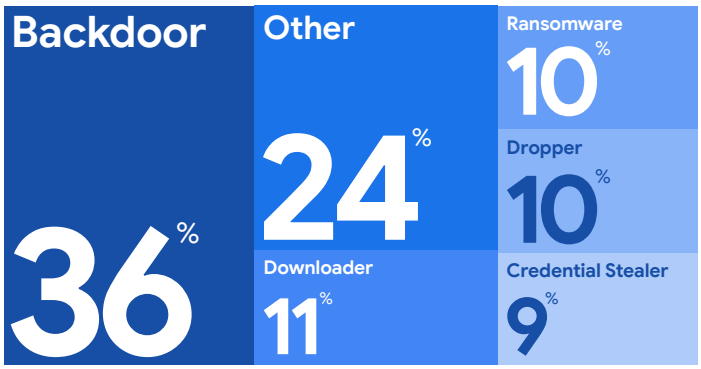

Malware Category	
Backdoor	A program whose primary purpose is to allow a threat cluster to interactively issue commands to the system on which it is installed.
Credential Stealer	A utility whose primary purpose is to access, copy, or steal authentication credentials.
Data Miner	A utility whose primary purpose is to gather ("mine") data, typically for theft by threat clusters. Excludes utilities that gather, 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 that 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 program whose primary purpose is to execute an external payload or shell command. A launcher does not contain or configure a payload it executes. Examples include a program that starts an executable file located on disk and a program that reads a payload from disk and executes it in memory.
Ransomware	A program whose primary purpose is to perform some malicious action (such as encrypting data) with the goal of extracting payment from the target 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 rootkits, keyloggers, and point-of-sale.

Most Frequently Seen Malware Families

The malware family most frequently observed in the investigations Mandiant performed in 2025 was GOLDVEIN.JAVA, the downloader used in a CL0P data theft extortion campaign exploiting zero-day CVE-2025-61882. The second-most frequently observed malware family was REDBIKE, also known as the Akira ransomware. After five consecutive years as the most frequently observed malware family in Mandiant investigations, Cobalt Strike BEACON fell to fourth.

Most Frequently Seen Malware Families, 2025

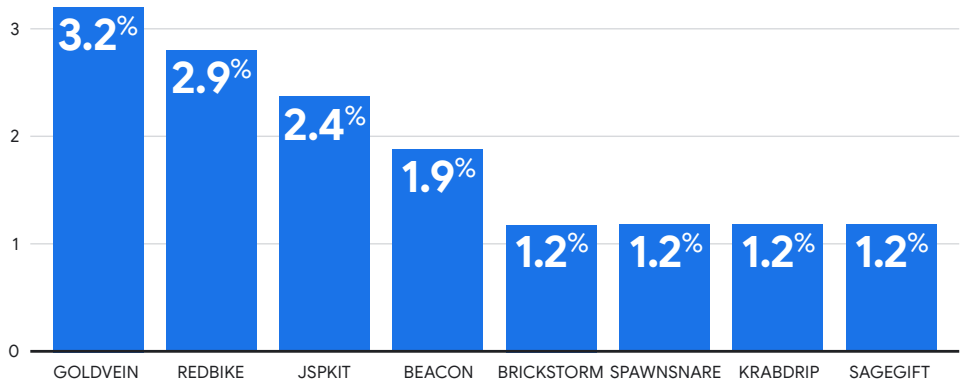

The proportion of BEACON malware observed in Mandiant investigations has declined each year since 2021, from 28% to just 2% in 2025.

BEACON Usage, 2020-2025

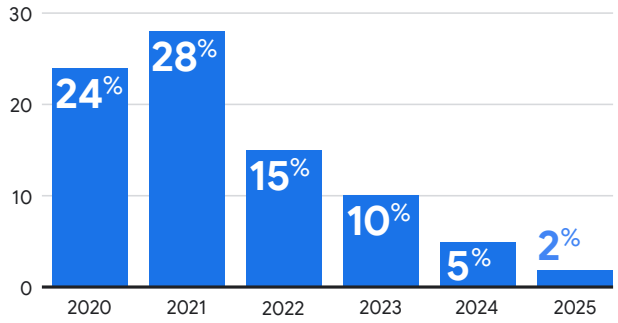

Malware Family	
GOLDVEIN.JAVA	GOLDVEIN.JAVA is a downloader written in Java for compromised Oracle WebLogic servers that connects to a hard-coded command-and-control (C2 or C&C) server over TCP to download an encrypted Java class payload. The downloaded payload is decrypted using a rolling XOR scheme with a hard-coded key and then directly invoked in memory through reflective loading. GOLDVEIN.JAVA is functionally similar to the original PowerShell version of GOLDVEIN; however, it uses the marker "TLSv3.1" in the C2 packet instead of the original "TLS v3".
REDBIKE	REDBIKE is a ransomware written in C++. This malware, also known as AKIRA, is designed to encrypt files on local drives and network shares. It operates as a stand-alone tool via command-line arguments that allow an operator to specify target paths, including files containing lists of network shares. REDBIKE employs a hybrid encryption scheme, using symmetric ciphers like ChaCha or KCipher-2 to encrypt file contents and an embedded RSA public key to protect the symmetric keys. More recent variants use a combination of ChaCha8 and KCipher-2 for file encryption. After encryption, it renames files by appending the ".akira" extension and creates a ransom note named "akira_readme.txt" in affected directories. To hinder recovery, the ransomware terminates a list of predefined processes and executes a PowerShell command to delete volume shadow copies. The malware avoids encrypting critical system files and directories by using an exclusion list for specific paths and extensions, such as ".exe," ".dll," and ".sys". REDBIKE does not contain a persistence mechanism and relies on an external tool or operator for execution. Some variants are verbose and will log their activity to a file in the user's temporary directory. REDBIKE exhibits some code overlap with CONTI ransomware.
JSPKIT	Publicly available JSP web shell known as JSP KIT or JSP SHELL or "cmd.jsp."
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.
BRICKSTORM	BRICKSTORM is a backdoor written in Go that communicates over WebSockets Secure (WSS). BRICKSTORM supports the ability to set itself up as a web server, perform file and directory manipulation, perform file operations such as upload/download, run shell commands, and perform SOCKS relaying.
SPAWNSNARE	SPAWNSNARE is a utility written in C that allows for extraction of the uncompressed Linux kernel image into a file and encrypts it using AES without the need for any other command-line utilities. Supported commands include writing a Bash script to disk, encrypting a file using AES-128 in CBC mode, decrypting a file using AES-128 in CBC mode, and acting as a BusyBox to run a set of statically linked applets. The dropped Bash script includes all of the necessary logic and commands for extraction of the kernel.
KRABDRIP	KRABDRIP is a downloader written in Rust. It uses an embedded configuration containing an AES key, IV, and an encrypted C2 URL to retrieve a payload. The malware communicates with its C2 server via HTTP to download an encrypted payload, which it decrypts using the embedded AES key and IV. The decrypted payload is treated as position-independent code and is injected into an instance of explorer.exe for execution. KRABDRIP performs anti-analysis checks, such as looking for a debugger and verifying the existence of a specific file in the temp directory. For evasion, it copies itself to the %TEMP% directory, relaunches from the new location, and then deletes both the original and copied files. The malware does not contain its own persistence mechanism.
SAGEGIFT	SAGEGIFT is an in-memory dropper written in Java for compromised Oracle WebLogic servers that reflectively loads an embedded Java class to execute in memory. SAGEGIFT returns its status through an encoded HTML comment in its response. SAGEGIFT has been observed loading SAGELEAF.

Threat Groups

What is an UNC group? When GTIG encounters new threat activity that cannot confidently be linked to an existing group, an UNC (or "uncategorized") 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, GTIG analysts build on the initial understanding of the attacker, potentially merging it with other tracked threat clusters and ultimately graduating the UNC to an advanced persistent threat (APT) or financial threat (FIN) group.

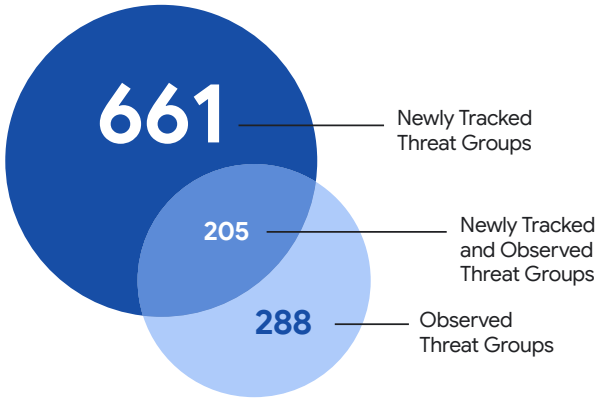

GTIG began tracking 661 net new threat clusters in 2025, bringing the total number of tracked threat clusters to more than 5,000. In 2025 investigations, Mandiant encountered 288 threat groups, 205 of which were newly tracked. The counts for these figures are slightly lower than the count of groups newly tracked and observed in 2024. While these figures align with a slight downward trend during the last five years, the consistent discovery of newly tracked groups underscores a decade-long trend of expanding threat activity. While "known" threats persist, the use of threat intelligence to continually identify and analyze new activity drives new detections and research, increasing defenders' ability to implement protections against both established and emerging adversaries.

Observed Groups by Goal

Financially motivated groups represent 41% of the threat clusters observed in the Mandiant investigations performed in 2025. This category includes six different North Korean threat groups seeking financial gain on behalf of the North Korean Government. Cyber espionage groups represented 16% of threat clusters in this data set. The proportion of cyber espionage groups observed in Mandiant investigations increased from 8% in 2024, while financially motivated groups declined from 55%. In 2025, Mandiant encountered threat clusters assessed to be motivated both by financial gain and a desire for notoriety; these clusters represented 4% of observed threat groups.

Observed Threat Groups by Goal, 2025

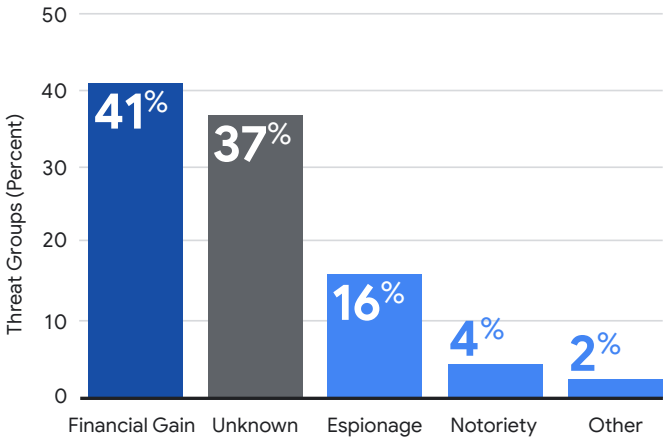

Ransomware

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

Extortion operations, including ransomware, data theft extortion, and multifaceted extortion, continue to represent the most impactful form of cyber crime. Ransomware-related intrusions accounted for 13% of Mandiant investigations in 2025. This is a result of the combined frequency of extortion incidents, and the disruption that extends beyond the targeted organization to affect clients, suppliers, and communities. Ransomware-related incidents Mandiant investigated in 2025 affected a wide range of organizations in the Americas; Europe, the Middle East, and Africa (EMEA); and Japan-Asia Pacific (JAPAC) including healthcare, technology, government, pharmaceutical, financial, education, aerospace and defense, business and professional services, and construction and engineering.

Initial Infection Vector

Across Mandiant-led ransomware investigations during 2025, we observed multiple ransomware operations relying on initial access partnerships, most commonly for malware distribution services, which we track as prior compromises. In 2025, prior compromise was the most frequently confirmed initial infection vector for ransomware-related incidents that Mandiant investigated. This is a marked increase from 15% observed in 2024, to 30% in 2025. In these incidents, the initial access threat cluster most commonly gained access via web compromises. GTIG tracks dozens of threat clusters that specialize in widespread initial access operations. These threat clusters focus on gaining an initial foothold at many organizations via high volume, opportunistic infection vectors, then selling or handing off this access to other threat clusters for post-compromise exploitation. This “hand-off” pattern is commonly observed in incidents that result in ransomware deployment.

The second-most common initial infection vector was exploits, at 27%. Brute-force attacks were the vector for 20% of ransomware-related intrusions, followed by stolen credentials and web compromise at 10% each.

Initial Infection Vector, 2025
Ransomware-Related

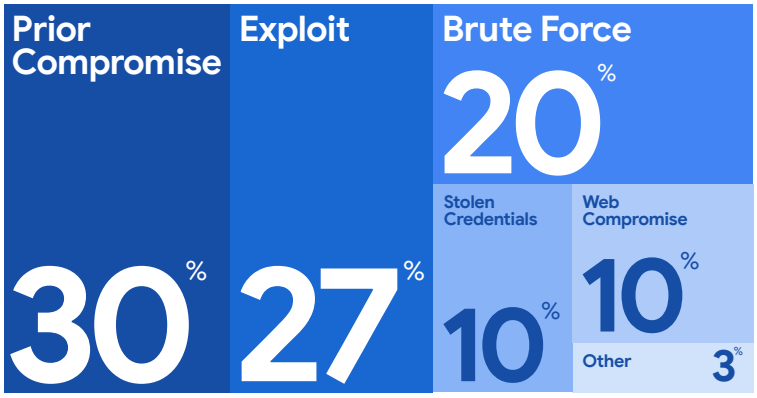

Detection by Source

Consistent with the extortion business model, in 2025, organizations most frequently learned of a ransomware incident from the attacker. In 44% of Mandiant's 2025 investigations, the intrusion was self-disclosed by the attacker either through direct extortion demands or the deployment of ransomware. Organizations discovered evidence of malicious activity internally in 41% of cases, and from an external entity such as law enforcement or a cybersecurity company in 15% of ransomware incidents. In 2024, adversary notifications comprised 49% of ransomware cases, and external entity notifications made up 21% of notification sources, figures roughly in line with 2025. However, internal discovery in 2024 represented 30%, compared to 41% in 2025.

Detection by Source, 2025

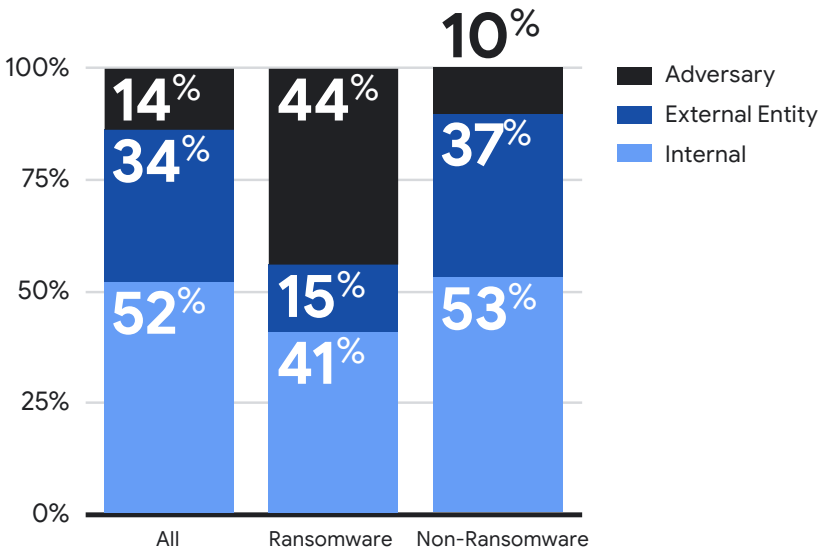

Ransomware vs. Global Dwell Time

Median dwell time for ransomware-related events in 2025 was nine days overall, five days for attacker-notified events, 12 days for internally discovered events, and eight days for compromises that organizations discovered due to notifications from external entities, such as law enforcement.

Dwell Time Distribution, Ransomware vs. Global

Comparing the distribution of dwell times for 2025 ransomware-related events to the overall dataset and non-ransomware-related incidents shows that ransomware-related incidents remain undiscovered for much shorter periods of time than other types of compromises. Ransomware-related intrusions were identified within one week 48.1% of the time, as opposed to 41.5% of the time for all incidents. In 2024, the concentration of ransomware-related dwell times in the one week or less category was slightly more pronounced, at 56.5%.

Global Dwell Time Distribution, 2025

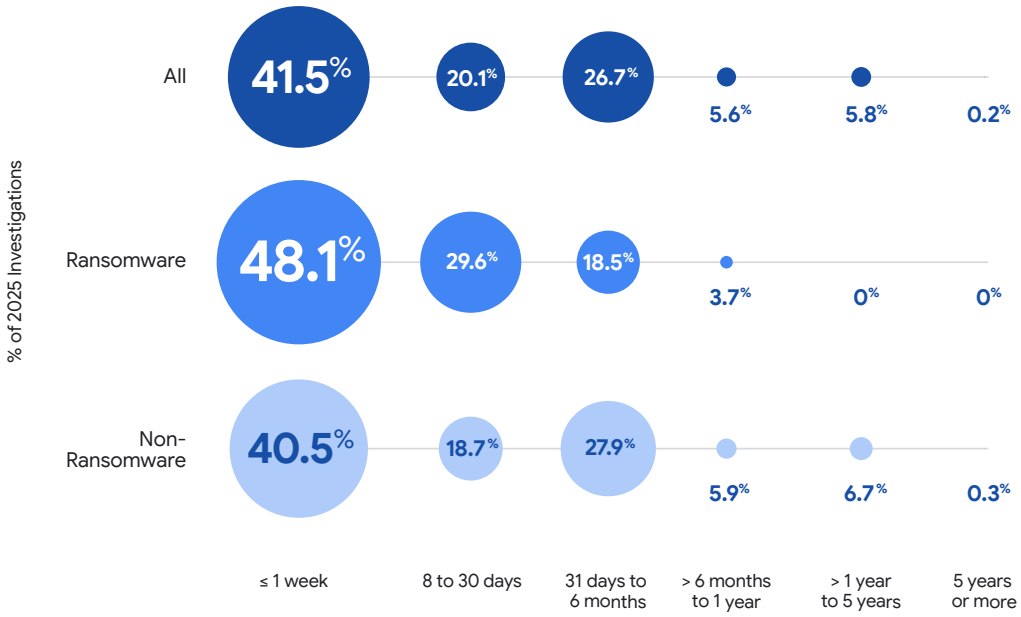

Malware

The top malware category for ransomware-related intrusions was, appropriately, ransomware (49%), followed by backdoors (19%), downloaders (10%), and tunnelers (7%). Compared to the global metrics for 2025, ransomware-related intrusions had less variety of malware categories. Two trends may contribute to this. First, ransomware operators often use legitimate utilities and remote monitoring and management (RMM) tools, which would not appear on this list. Second, the ransomware-related dataset does not include initial access operations prior to the hand-off, which would be characterized by more frequent use of droppers, launchers, and credential stealers.

Observed Malware Families by Category, 2025
Ransomware-Related

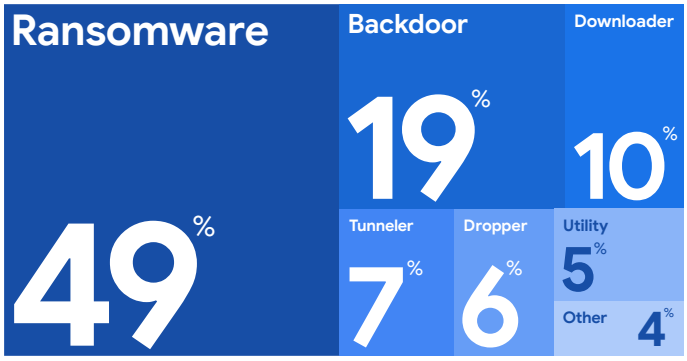

The most frequently observed malware family in 2025 Mandiant ransomware-related incident response investigations was REDBIKE (aka Akira) at 15%, followed by AGENDA (aka Qilin) and ADAPTAGENT, the backdoor component of the publicly available AdaptixC2 pentesting framework, both at 5%. The INC ransomware variant was also in the top-observed malware category, observed in 4% of ransomware-related investigations.

The SYSTEMBC tunneler was the fifth-most frequently seen malware family in 2025 ransomware-related intrusions, reflecting its popularity among multiple groups that conduct ransomware operations. SYSTEMBC was among the most frequently observed malware families in 2024 ransomware-related instructions as well. CHERRYQUIRK has been observed downloading and executing ADAPTAGENT.

Most Frequently Seen Malware, 2025
Ransomware-Related

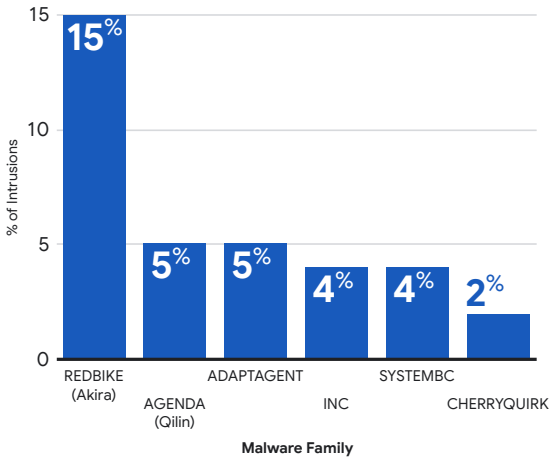

Mandiant observed commercially available and legitimate tools used much more frequently than malware in ransomware-related intrusions in 2025. The most common category of tools was utilities, which make up 35% of the dataset. The second-most frequently seen type of tools were reconnaissance tools, followed by archivers, such as WinRAR, remote control and administration tools, lateral movement tools, and credential stealers such as Mimikatz.

Observed Tools, 2025
Ransomware-Related

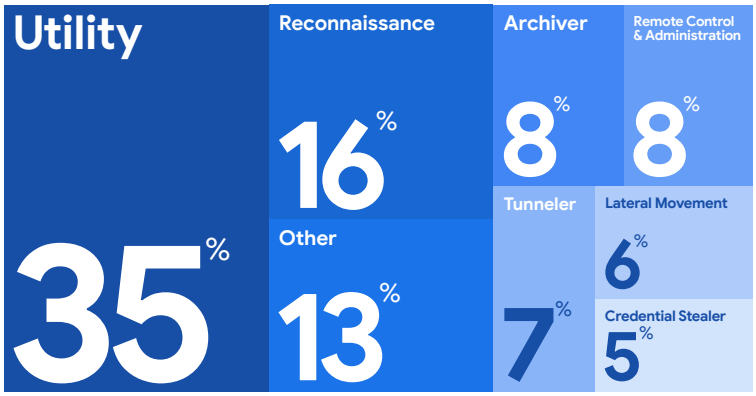

Tool Roles

Utility	A utility is a program that has a specialized purpose that does not fit into any other defined category (such as keylogger, sniffer, or credential theft). Examples may include tools designed to overwrite or clear log files, encode or decode files, etc.
Reconnaissance Tool	A reconnaissance tool is 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).
Archiver	An archiver is a program whose primary purpose is to package one or more files into an archive and may also extract files from an existing archive. The program may have additional options to compress or encrypt the archived files. Common examples include RAR, ZIP, and TAR.
Remote Control and Administration Tool	A remote control and administration tool is a legitimate program whose primary purpose is to remotely access and control or administer a system.
Tunneler	A tunneler is a program that proxies or tunnels network traffic.
Lateral Movement	A lateral movement tool is a program whose primary purpose is to facilitate lateral movement within a network.
Credential Stealer	A credential stealer is a utility whose primary purpose is to access, copy, or steal authentication credentials.

The tools Mandiant encountered most often in 2025 ransomware-related events were the RAR archiver, which is commonly used to prepare data for theft, and nltest, which is a utility for interacting with domain controllers and domain services. Ransomware operators frequently use nltest to conduct reconnaissance or within ransomware deployment scripts. Threat clusters take advantage of the useful functionalities of all of these tools to accomplish their malicious activities as well as reduce the likelihood that their activity will be detected as malicious.

Most Frequently Seen Tools, 2025
Ransomware-Related

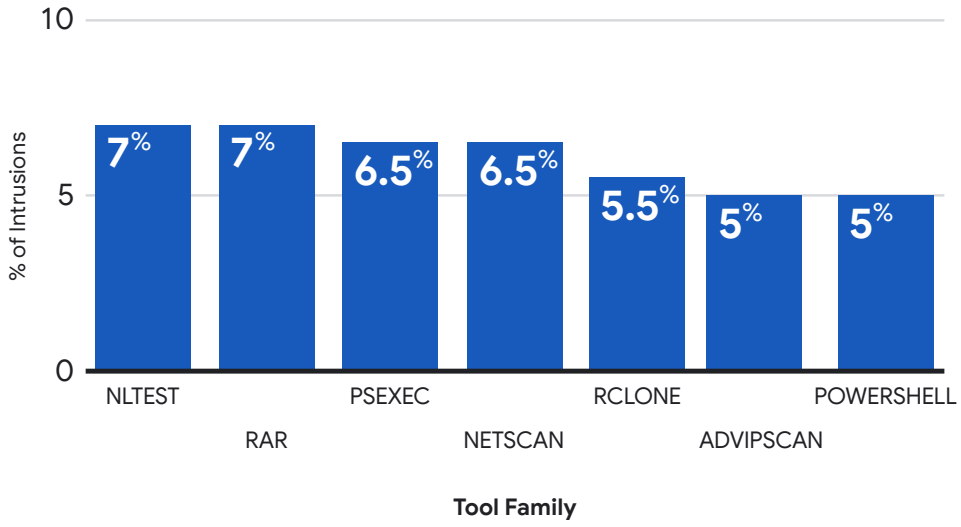

Tools	
Nltest	Nltest is a Windows command-line utility used to list domain controllers and information about domain trust relationships.
RAR	RAR is the RAR command-line archive utility; it may be distributed as part of the WinRAR package.
PsExec	PsExec, a utility within the Microsoft Sysinternals suite created by Mark Russinovich, is a command-line tool that executes processes on remote Windows systems. While designed for legitimate remote administration, it is frequently used by threat clusters for lateral movement within compromised networks because it does not require pre-installing software on the target machine.
SoftPerfect Network Scanner	NetScan is a SoftPerfect network scanner, 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 Windows Management Instrumentation (WMI), Simple Network Management Protocol (SNMP), HTTP, and NetBIOS.
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, File Transfer Protocol (FTP), Google Drive, HTTP, Mega, Microsoft OneDrive, rsync.net, Secure File Transfer Protocol (SFTP), and the local file system.
Advanced IP Scanner	Advanced IP Scanner (ADVIPSCAN) is a publicly available network scanner developed by Famatech that has remote control capabilities.
PowerShell	PowerShell is a publicly available cross-platform task automation and configuration management utility framework, consisting of a command-line shell and scripting language. Unlike most shells, which accept and return text, PowerShell is built on top of the .NET Common Language Runtime (CLR) and accepts and returns .NET objects.
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.
FileZilla	FileZilla is a cross-platform, publicly available FTP utility.
Impacket	Impacket is a Python library that allows users to work with various network protocols.
AnyDesk	AnyDesk is a commercially available remote monitoring and management (RMM) application that is supported on Windows, macOS, Linux, Android, and ChromeOS devices.

Ransomware Operations

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 ransomware-as-a-service (RaaS) business model is characterized by threat clusters specializing in components of the threat lifecycle, such as initial access operations, and interoperability between threat groups and ransomware brands. A centralized group may maintain the ransomware itself as well as the data leak site (DLS) infrastructure, while numerous affiliates deploy the ransomware. While the total volume of DLS posts has increased in 2025, disruptions, such as law enforcement operations or internal conflict between clusters, have led to the disappearance or decline of previously prolific RaaS groups like LockBit and ALPHV (BlackCat). However, the well-established AGENDA (Qilin) and REDBIKE (Akira) RaaS brands emerged as the most prolific DLS in 2025.

2025 DLS Listings for Qilin and Akira

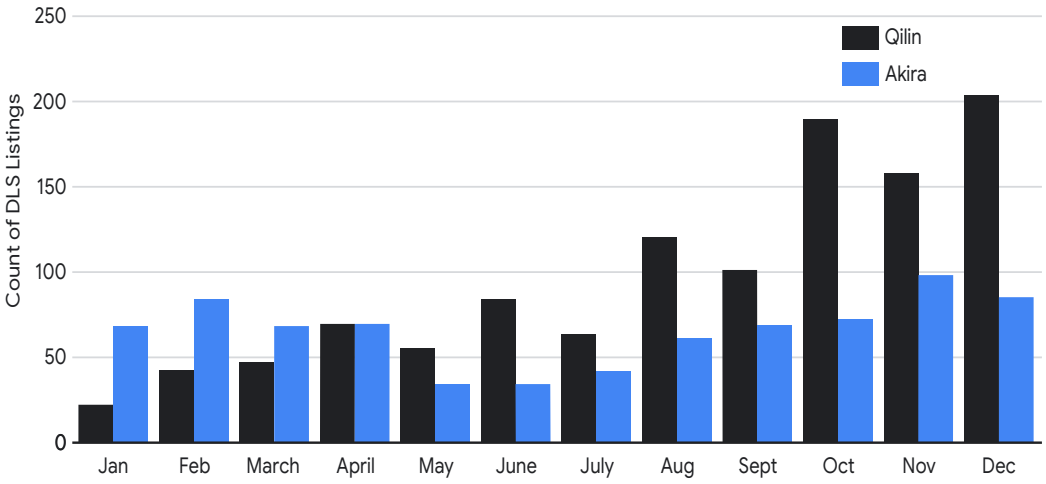

AGENDA (aka Qilin)

The Qilin RaaS and associated DLS first emerged in 2022. In 2025, it became the single-most prolific RaaS based on count of DLS listings, though AGENDA was the second-most commonly observed ransomware in Mandiant incident response in 2025. GTIG tracks multiple threat clusters that have deployed AGENDA, which is the name GTIG uses to track the ransomware commonly referred to as Qilin, including UNC2465 and UNC6276.

UNC2465 is a financially motivated threat cluster that has been active since at least April 2019. They frequently obtain access to environments via malicious installers, masquerading as legitimate software—usually IT administration tools—that lead to the SMOKEDHAM backdoor. SMOKEDHAM may be exclusive to UNC2465. The threat cluster has used malicious advertising (malvertising) for malware distribution, but has often sought traffic providers as well. UNC2465 has remained interested in monetizing accesses via ransomware operations and is likely leveraging AGENDA ransomware in their current operations; they have previously used HIVELOCKERS.HUNTERS (aka Hunters International), LOCKBIT, and DARKSIDE ransomware.

UNC2465 SMOKEDHAM LOCKBIT CAMP.25.017

UNC6276 is a financially motivated threat cluster that has been active since at least May 2025 and has targeted organizations in North America, including legal and professional services and manufacturing, to deploy ransomware. GTIG observed UNC6276 gain initial network access through the use of either stolen or brute-forced virtual private network (VPN) credentials. UNC6276 has leveraged publicly available tools such as PINGCASTLE, NETSCAN, and MIMIKATZ for reconnaissance and credential theft and deployed SYSTEMBC.LINUX for command and control. These intrusions have led to the deployment of AGENDA ransomware and data theft extortion.

UNC6276 SYSTEMBC.LINUX CAMP.25.058

REDBIKE (aka Akira)

The Akira RaaS and DLS were established in 2023, and the RaaS has increased its market share, as measured by count of DLS listings, steadily since then to become the second-most prolific DLS in 2025. REDBIKE, which is the name GTIG uses to track the ransomware, was the second-most frequently seen malware family across all Mandiant investigations in 2025, and the most commonly observed ransomware variant. GTIG tracks multiple threat clusters using REDBIKE, including UNC6361 and a suspected FIN6 cluster.

UNC6361 compromises target environments through the exploitation of known vulnerabilities in network edge devices. These compromises lead to lateral movement further into the target environment, followed by the eventual deployment of REDBIKE ransomware in support of this cluster's extortion operations.

UNC6361 REDBIKE CAMP.25.051

FIN6 is a financially motivated threat cluster active since mid-2014. Since mid-2019, FIN6 has used job-themed lures and fake personal websites to deliver BULLZLINK, followed by SQUIDSLEEP and SQUIDGATE. Notably, in 2025 FIN6 used a financial-themed lure to deliver an updated version of SQUIDGATE. While GTIG has not directly observed FIN6 monetize access recently, overlaps identified with suspected affiliates suggest the group likely supports or conducts REDBIKE ransomware operations. Historically, the group compromised point-of-sale (POS) environments using TRINITY (aka FrameworkPOS), used SCRAPMINT malware to steal payment card data, and targeted card-not-present (CNP) data in e-commerce environments. As of mid-2018, at least one FIN6-affiliated cluster began deploying ransomware, including LOCKERGOGA, RYUK, MEGACORTEX, and MAZE.

Cloud Compromises

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

For more **information about UNC6240**, please see:

Vishing for Access: Tracking the Expansion of ShinyHunters-Branded SaaS Data Theft

For more **information about BRICKSTORM** activity, please see:

Another BRICKSTORM: Stealthy Backdoor Enabling Espionage into Tech and Legal Sectors

For the **VMware vSphere hardening guide**, please see:

From Help Desk to Hypervisor: Defending Your VMware vSphere Estate from UNC3944

In 2025, the most common initial infection vector found during Mandiant investigations of cloud-related compromises was voice phishing, at 23%, followed by third-party compromise (17%), stolen credentials (16%), email phishing (15%), insider threat (14%), and exploits (6%). Mandiant identified evidence of data theft in 59% of cloud compromises. Just over a third of cases, 34%, supported financially motivated objectives, including employment fraud, data theft extortion, ransomware, payment redirection fraud, and theft.

Voice phishing facilitated significant data theft extortion campaigns attributed to UNC3944 and UNC6240 in 2025. In terms of post-compromise operations in cloud environments, UNC3944 conducted extensive reconnaissance of targeted organizations' cloud resources, including SharePoint, Azure Portal, M365 email, and privileged access management (PAM) solutions, and extracted sensitive data for attempted extortion. UNC6040 used native utilities to automate large-scale data collection and theft. In several cases, UNC6040 used bulk application programming interface (API) operations to extract broad datasets to use for extortion.

Cloud Initial Infection Vectors, 2025

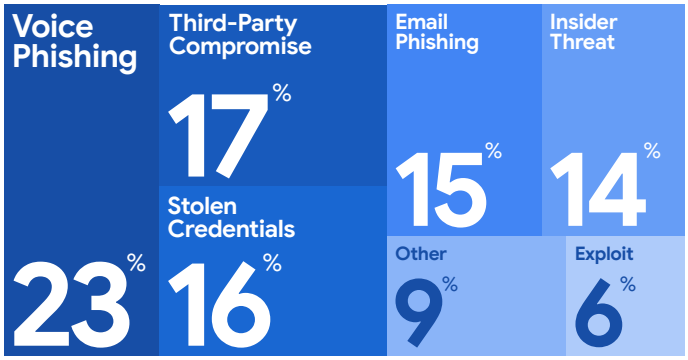

In another significant set of activity targeting cloud environments, GTIG tracked the PRC-nexus cyber espionage cluster UNC6201's use of stealthy tactics and lightweight malware to maintain long-term access to targeted environments. Mandiant investigated a number of incidents in 2025 in which the threat cluster deployed the BRICKSTORM backdoor on appliances that do not support endpoint detection and response (EDR), including Linux- and BSD-based appliances from multiple manufacturers. Using valid credentials likely captured on the network device, UNC6201 then accessed VMware vCenter servers and ESXi hosts. With access to vCenter, the threat cluster cloned virtual machines (VMs), which included single sign-on (SSO) identity providers, secret vaults, and domain controllers. By accessing targeted data and credentials in the cloned but powered off VMs, the threat cluster circumvented security alerting on those systems. Mandiant also identified evidence of UNC6201 leveraging this access to target and datamine sensitive cloud-based resources, such as mailboxes belonging to developers and system administrators as well as individuals involved in matters that align with the economic and espionage interests of the People's Republic of China.

Artificial Intelligence

GTIG has closely tracked threat cluster interest in, as well as use and misuse of, artificial intelligence (AI) in malicious operations. In 2025, threat clusters have increasingly adopted AI tools to achieve productivity gains in different stages of the attack lifecycle, particularly in tasks such as reconnaissance, social engineering, and malware development. In 2025, Mandiant investigations identified threat clusters using AI-themed lures, stealing credentials for AI applications, and targeting companies developing AI technologies.

Notably, threat clusters have also relied on AI tools within the compromised environment to help carry out their operations. For example, Mandiant investigated an NPM package manager software supply chain compromise that led to the installation of the QUIETVAULT credential stealer. Upon activation, QUIETVAULT checks to see if AI command-line interface (CLI) tools are installed on the targeted machine, and if so, executes a predefined prompt to search for configuration files. The tool then attempts to collect GitHub and NPM tokens and, if found, copy them to a publicly accessible GitHub repository.

Threat Actor Use and Abuse of AI

For more information about threat actor use and abuse of AI, please see:

GTIG AI Threat Tracker: Distillation, Experimentation, and (Continued) Integration of AI for Adversarial Use

Adversarial Misuse of Generative AI

GTIG AI Threat Tracker: Advances in Threat Actor Usage of AI Tools

UNC1069 Targets Cryptocurrency Sector with New Tooling and AI-Enabled Social Engineering

Text-to-Malware: How Cybercriminals Weaponize Fake AI-Themed Websites

AI Risk and Resilience: A Mandiant Special Report

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 2025, MITRE released ATT&CK version 18, a release that updates Techniques, Groups, Campaigns and Software for Enterprise, Mobile, and ICS. Similar to version 16, this change did not introduce a significant number of new techniques and sub-techniques to the already established framework. 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

Top 10 Most Frequently Seen MITRE ATT&CK Techniques

Rank	Technique	Percent
1	T1059: Command and Scripting Interpreter	45.9%
2	T1074: Data Staged	39.6%
3	T1083: File and Directory Discovery	33.5%
4	T1021: Remote Services	30.6%
5	T1190: Exploit Public-Facing Application	27.7%
6	T1027: Obfuscated Files or Information	27.2%
7	T1070: Indicator Removal	24.8%
8	T1105: Ingress Tool Transfer	24.5%
9	T1033: System Owner/User Discovery	24.0%
10	T1133: External Remote Services	23.8%

MITRE ATT&CK Sub-Techniques Used Most Frequently

Top 10 Most Frequently Seen MITRE ATT&CK Sub-Techniques

Rank	Technique	Percent
1	T1059.003: Windows Command Shell	26.2%
2	T1059.001: PowerShell	24.0%
3	T1021.001: Remote Desktop Protocol	22.1%
4	T1021.002: SMB/Windows Admin Shares	19.4%
5	T1204.002: Malicious File	16.5%
6	T1070.004: File Deletion	15.5%
7	T1505.003: Web Shell	14.8%
8	T1569.002: Service Execution	14.3%
9	T1016.001: Internet Connection Discovery	14.1%
10	T1021.004: SSH	12.9%

Command and Scripting Interpreter (T1059) remains a top MITRE ATT&CK tactic, technique, and procedure (TTP) for the fifth consecutive year, followed by Data Staged (T1074), File and Directory Discovery (T1083), and Remote Services (T1021). The consistency in Command and Scripting Interpreter (T1059) is expected given, by MITRE's own definition, most systems come with some built-in command-line interface and scripting capabilities. These built-in command-line capabilities allow threat clusters to leverage living off the land (LotL) to perform actions across multiple stages of the attacker lifecycle, with the added benefits of convenience and lowered chances of getting detected. The latter is also why Indicator Removal (T1070) remains a consistent top observed TTP year over year. Remote Services (T1021), System Owner/User Discovery (T1033), and File and Directory Discovery (T1083) are foundational to post-exploitation internal reconnaissance and privilege escalation. As denoted in the Mandiant Attack Lifecycle, threat clusters will often perform internal reconnaissance and move laterally after gaining access to environments, whether to escalate privileges, maintain access, or quickly gather information and steal data, potentially to be used later in extortion.

Several techniques returned to the Top 10 after a multi-year hiatus, including Data Staged (T1074), File and Directory Discovery (T1083), and Ingress Tool Transfer (T1105). This resurgence, paired with a similarly observed increase in System Owner/User Discovery (T1033), shows heavy emphasis on internal reconnaissance and lateral movement. Notably, Data Encrypted for Impact (T1486) was not in the top 10, where it appeared for the first time in 2024. While this indicates a decline in encryption-based ransomware operations, the rise of Data Staging (T1074) suggests a tactical shift toward data theft and pure extortion-based models.

Regional Breakouts

Americas

The metrics reported in this section are based on Mandiant investigations affecting organizations in North, Central, and South America.

AMERICAS

Exploit
28%

Voice Phishing
14%

Web Compromise
10%

Initial Infection Vector

In 2025, exploits (28%) remained as the leading entry for compromises in the Americas where Mandiant was able to determine an initial infection vector. While email phishing saw a significant decline—down to 5% from 16% in 2024—web compromises claimed a spot in the top three. This follows similar trends observed globally, where similar declines in email phishing and increases in web compromises were noted.

Detection by Source

In 2025, GTIG observed an increase in external notifications for activity Mandiant investigated in the Americas. This increase brings external notification in line with internal notifications in an even distribution. Of these external notifications, one-third (33%) originated from external partners such as law enforcement and cybersecurity companies, while the remaining 17% came directly from threat clusters via ransom notes or extortion attempts.

Americas Detection by Source, 2017-2025

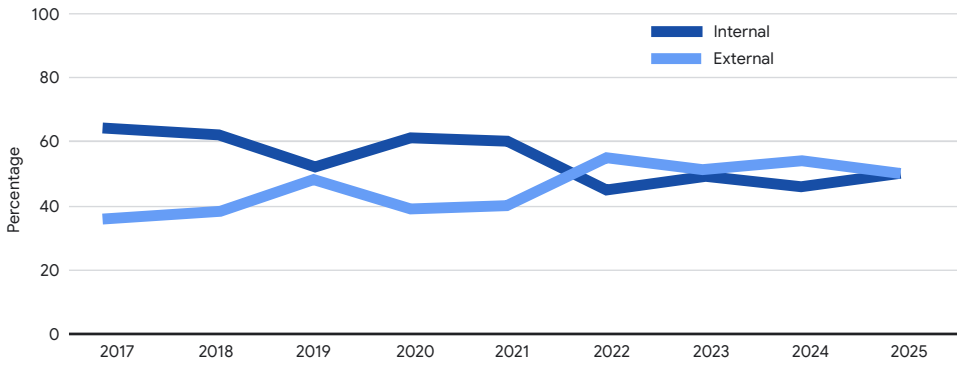

Median Dwell Time

This year, the median dwell time for intrusions Mandiant investigated in the Americas in 2025 was 12 days overall, an increase of two days. The median dwell time for internally notified events was nine days, consistent with last year's median dwell time. However, externally notified events saw an increase to 17 days, which is a notable increase from last year's 10-day median dwell time.

Americas Median Dwell Time, 2016-2025

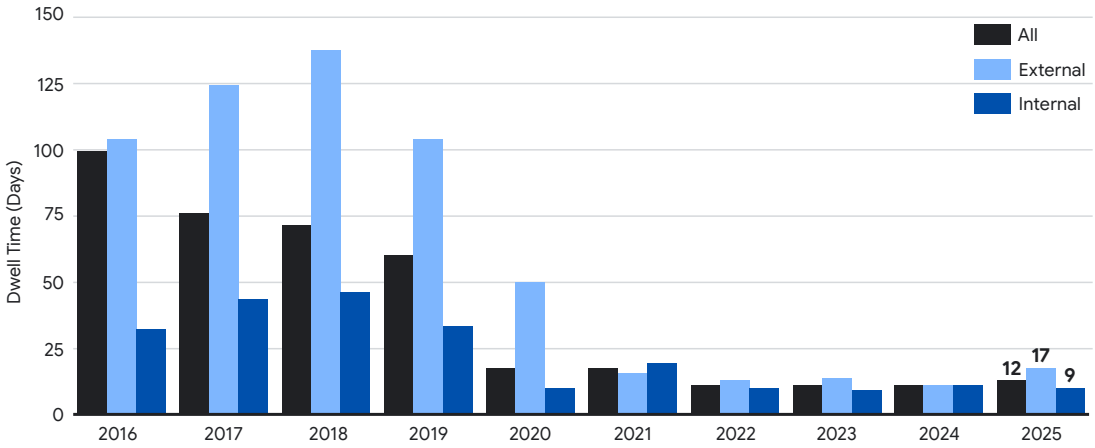

The 2025 dwell time distribution for the Americas also shows increases in longer duration incidents compared to 2024, specifically, incidents lasting between one and six months, and one to five years. Similar to longer dwell times observed globally, this shift in the Americas likely reflects a larger proportion of incidents in which threat actors prioritize maintaining long-term access to targeted environments, including cyber espionage, North Korean IT workers, and other types of compromises.

Americas Dwell Time Distribution, 2021-2025

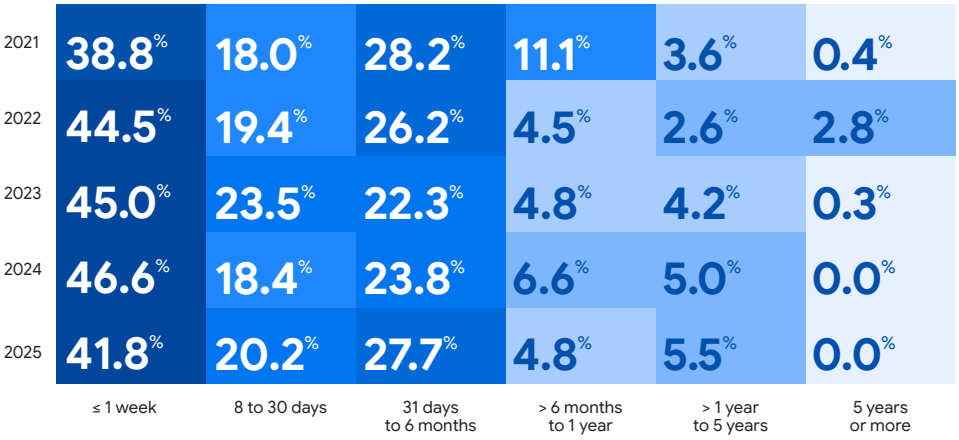

Threat Groups

One of the most frequently observed threat clusters in Mandiant engagements in the Americas in 2025 was UNC5267, which is the designation GTIG uses to track most North Korean IT worker activity. This set of activity likely contributed to the observation that greater proportions of incidents in the Americas had longer dwell times in 2025, as North Korean IT workers seek to remain undetected and collect paychecks as long as possible.

Another frequently observed threat cluster was UNC6395. This threat cluster compromised a software-as-a-service (SaaS) provider and stole authentication tokens for many of the organization's clients. Then UNC6395 used those tokens to access the client environments, download code repositories, and identify additional credentials and keys contained within them.

Campaigns and Global Events

In 2025, the GTIG tracked multiple campaigns and global events that directly correlated to Mandiant investigations in the Americas. The Campaigns and Global Events timeline summarizes prominent threat activity observed in the region.

Please see the graphic on the following page for more detail.

2025 Americas Campaigns and Global Events Related to Mandiant Incident Investigations

EMEA

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

Initial Infection Vector

For Mandiant investigations affecting EMEA in 2025, the top initial infection vectors were exploits (51%), email phishing (12%), and prior compromise (9%). Exploits and prior compromise were the first- and third-most commonly observed vectors globally in 2025, but email phishing had a higher prevalence in EMEA than in global statistics. In 2024 investigations, the top two categories were also exploits and email phishing, though the third-most common vector was brute-force attacks.

EMEA

Exploit
51%

Email Phishing
12%

Prior Compromise
9%

Detection by Source

Organizations first detected evidence of malicious activity internally in 60% of the compromises Mandiant investigated in EMEA in 2025, and external notifications were the first indication in 40% of cases. This is a reversal from 2024, when internally discovered incidents represented 41%, and externally notified events represented 59%

EMEA Detection by Source, 2017-2025

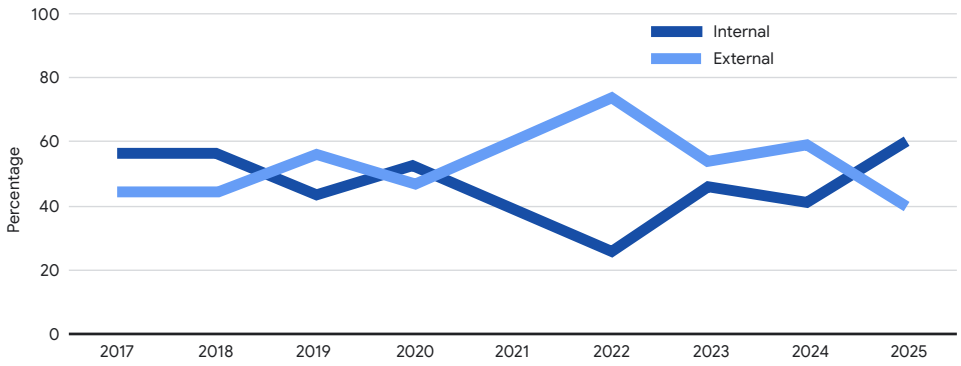

Median Dwell Time

The median dwell time for Mandiant investigations in EMEA in 2025 was 20 days, 19 days for internally discovered incidents, and 21 days for externally notified events. Compared to 2024, the overall median dwell time for incidents investigated decreased by seven days.

EMEA Median Dwell Time, 2016-2025

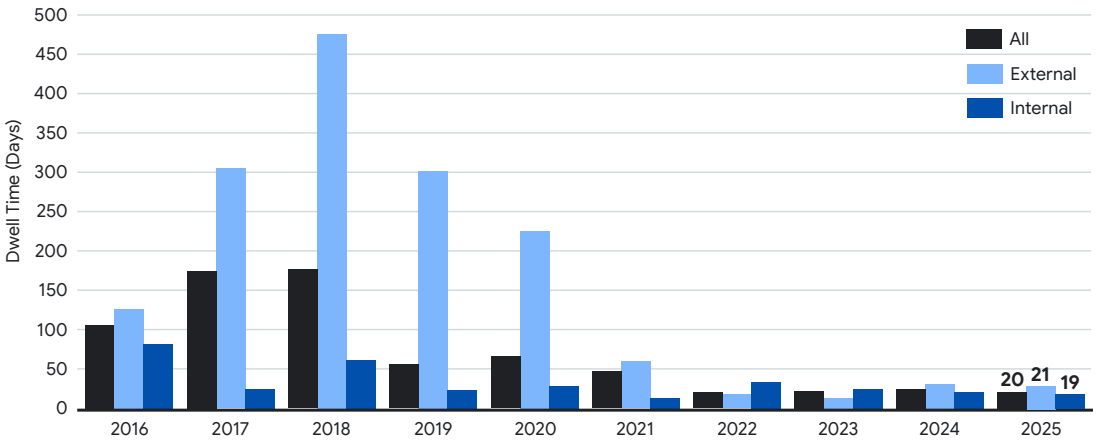

The EMEA dwell time distribution reveals that over the long term, dwell times for incidents Mandiant investigated in EMEA have grown more concentrated towards shorter dwell times. For example, incidents discovered within one week increased from 36.7% in 2024 to 40.3% in 2025, and incidents discovered between one week and one month increased from 16.5% to 19.4%.

EMEA Dwell Time Distribution, 2021-2025

	≤ 1 week	8 to 30 days	31 days to 6 months	> 6 months to 1 year	> 1 year to 5 years	5 years or more
2021	33.0%	14.0%	22.0%	12.0%	14.0%	6.0%
2022	41.6%	12.2%	17.7%	10.2%	11.5%	7.0%
2023	35.9%	20.5%	23.1%	6.4%	14.1%	0.0%
2024	36.7%	16.5%	27.8%	3.8%	12.7%	2.5%
2025	40.3%	19.4%	25.4%	4.5%	9.0%	1.5%

For more **information about UNC1549**, please see:

Frontline Intelligence: Analysis of UNC1549 TTPs, Custom Tools, and Malware Targeting the Aerospace and Defense Ecosystem

Threat Groups

Mandiant investigated several incidents affecting organizations in EMEA related to two distinct activity clusters exploiting CVE-2025-31324 as a zero-day. Observed activity was largely limited to establishing an initial foothold, for example through dropping the JSPKIT web shell.

In 2025, Mandiant encountered multiple compromises that GTIG attributes to the Iranian cyber espionage threat cluster UNC1549 in EMEA. Prior to discovery, the threat cluster had maintained access to compromised organizations for periods spanning months to upwards of two years. UNC1549 deployed custom malware including MINIBIKE and TWOSTROKE backdoors and targeted credentials both for privilege escalation and for potential utility in targeting additional organizations.

Campaigns and Global Events

In 2025, the GTIG tracked multiple campaigns and global events that directly correlated to Mandiant investigations in EMEA. The Campaigns and Global Events timeline summarizes prominent threat activity observed in the region.

Please see the graphic on the following page for more detail.

2025 EMEA Campaigns and Global Events Related to Mandiant Incident Investigations

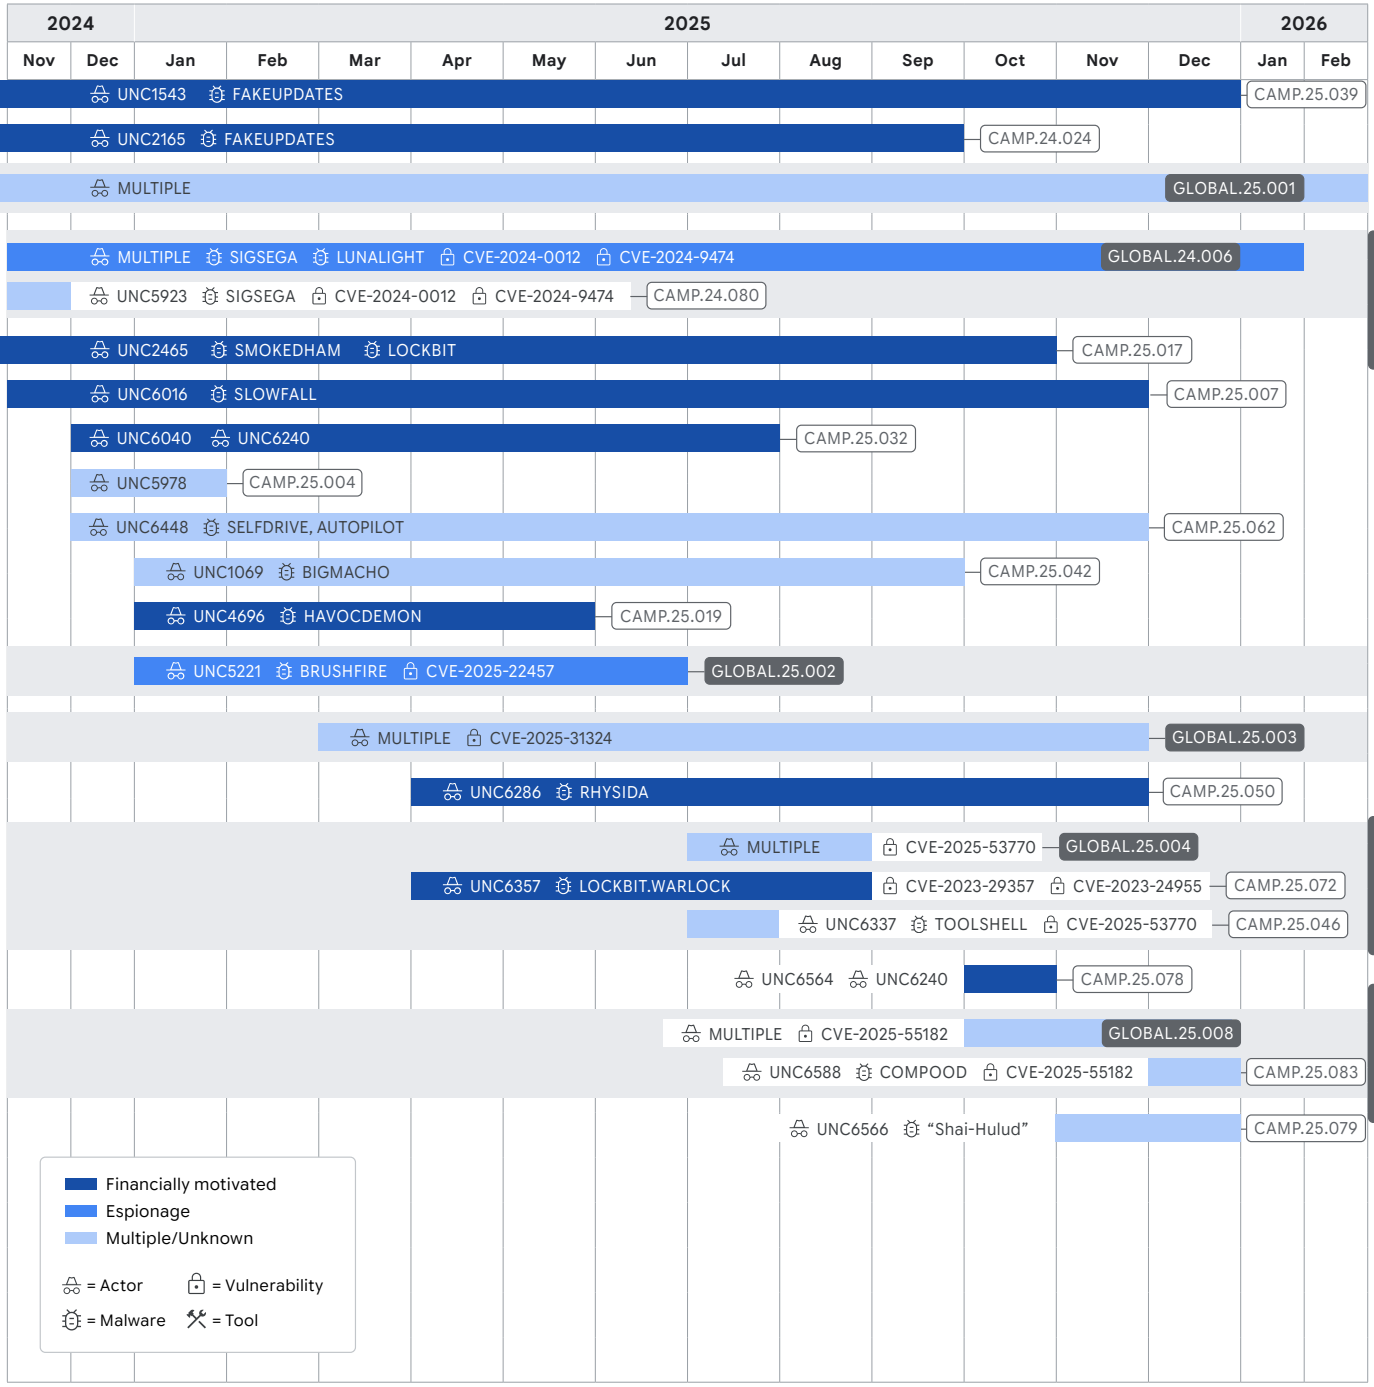

JAPAC

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

Initial Infection Vector

The most frequently observed initial infection vector for 2025 Mandiant investigations in JAPAC was exploits, at 33%, followed by stolen credentials at 23%. The ranking of these two vectors is consistent with 2024 investigations; however, exploits were much more prominent in 2024, representing the vector for 64% of compromises. The third-most common vector in 2025 investigations in JAPAC was prior compromise at 16%.

JAPAC

Exploit
33%

Stolen Credentials
23%

Prior Compromise
16%

Detection by Source

Organizations were notified of a potential compromise from an external organization in 51% of the compromises Mandiant investigated in JAPAC in 2025, while internally discovered incidents represented 49% of the dataset. This is a significant shift compared to 2024, when external notifications accounted for 69% of investigations, and internal discovery 31%.

JAPAC Detection by Source, 2017-2025

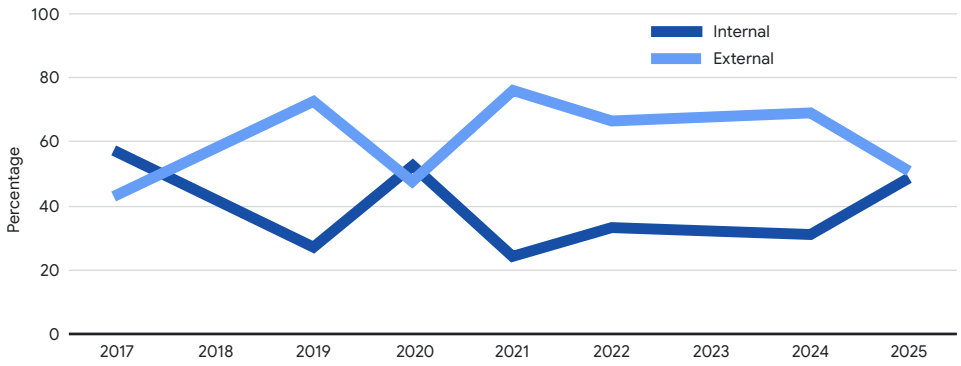

Median Dwell Time

Median dwell time for compromises Mandiant investigated in JAPAC in 2025 was 15 days overall, seven days for internally discovered incidents, and 38 days for externally notified incidents. The overall and externally notified medians are significantly higher than observed in 2024—6 and 10, respectively. The increases in 2025 likely reflect the presence of a number of intrusions in which threat clusters prioritized stealth and remained undiscovered for extended periods, including cyber espionage incidents as well as financially motivated clusters conducting reconnaissance for months before executing payment transfer fraud.

JAPAC Median Dwell Time, 2016-2025

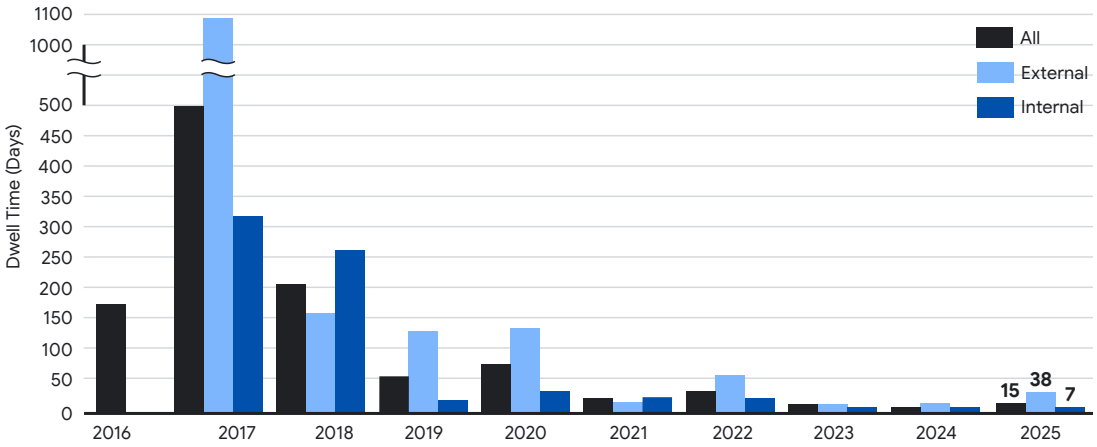

The dwell time distribution for JAPAC investigations also reflects this shift, with greater proportions of 2025 incidents falling into the one week to a month (21.6%), one month to six months (21.6%) and six months to one year (11.8%) time windows compared to 2024.

JAPAC Dwell Time Distribution, 2021-2025

Threat Groups

In 2025, multiple Mandiant investigations in JAPAC were attributed to a suspected PRC-nexus cyber espionage cluster exploiting CVE-2025-31324 after the patch date in April 2025. This threat cluster sent basic reconnaissance commands to web shells installed on vulnerable, compromised devices, but it is unclear whether this cluster installed the web shells.

Mandiant encountered another threat cluster in multiple JAPAC compromises in 2025 that also exploited CVE-2025-31324 after the patch date. This threat cluster typically installed the KRABDRIP downloader and conducted reconnaissance.

Campaigns and Global Events

In 2025, GTIG tracked multiple campaigns and global events that directly correlated to Mandiant investigations in the JAPAC region. The Campaigns and Global Events timeline summarizes prominent threat activity observed in the region.

Please see the graphic on the following page for more detail.

2025 JAPAC Campaigns and Global Events Related to Mandiant Incident Investigations

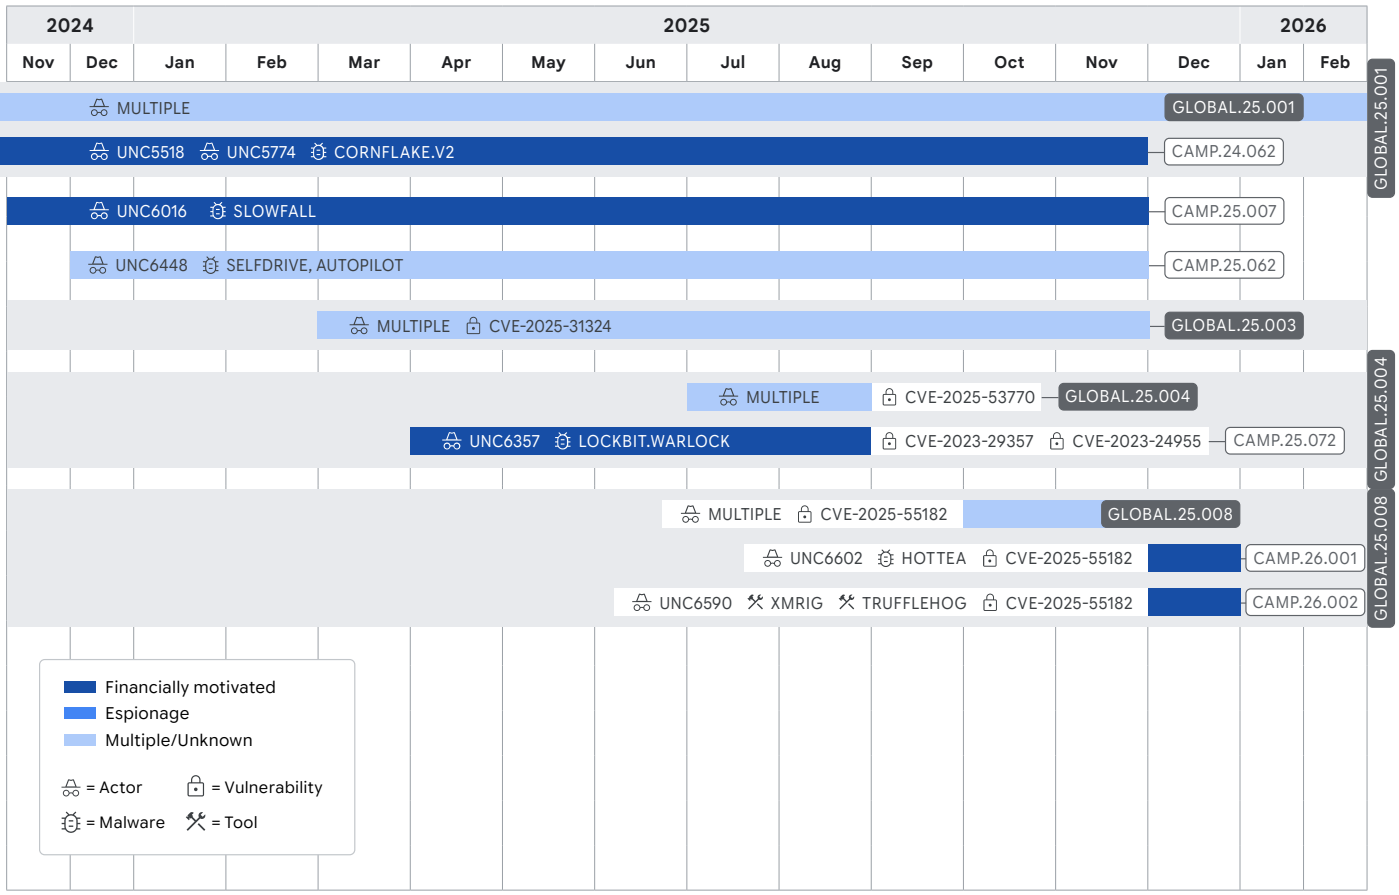

Articles

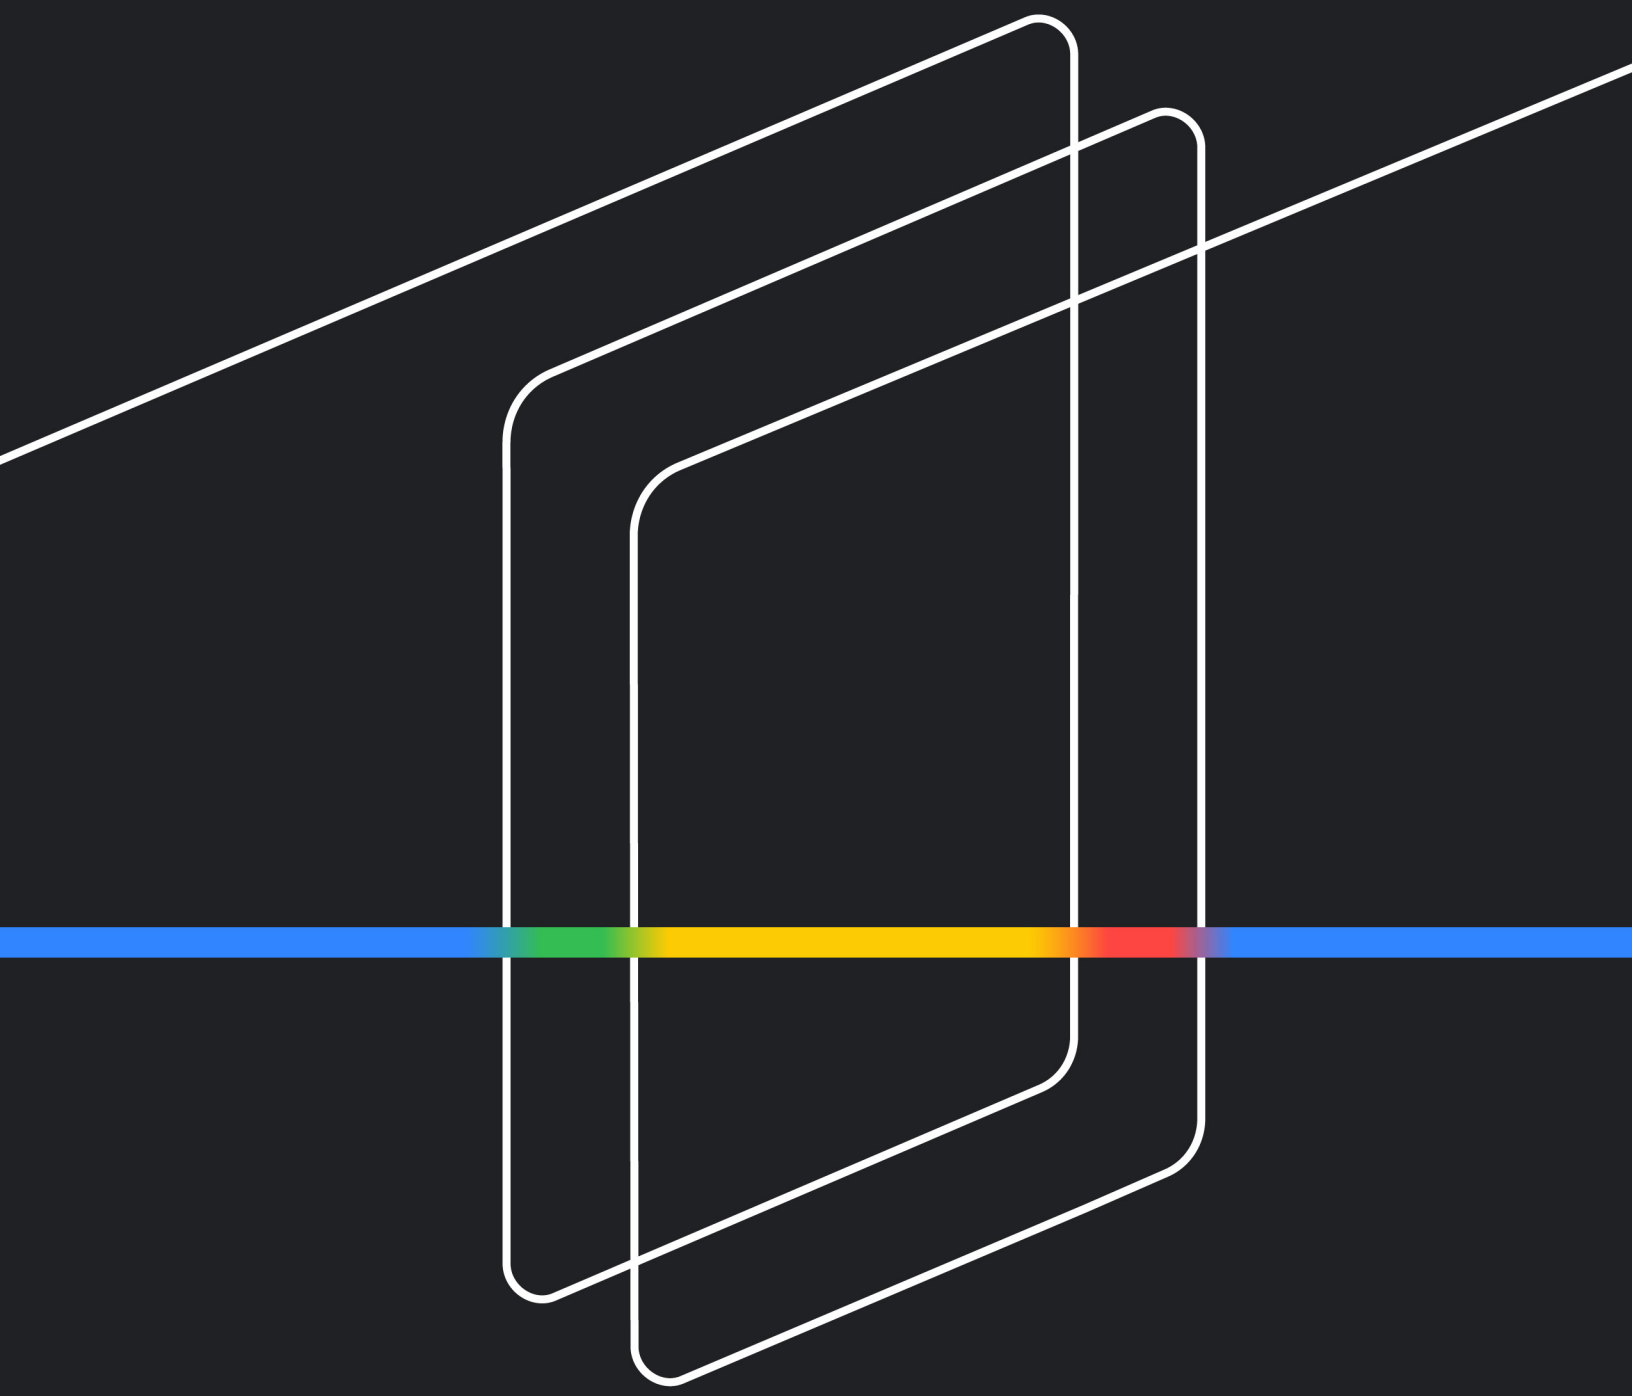

A Minor Infection Today Can Be a Ransomware Attack Tomorrow

In recent years, Mandiant has observed a gradual upward trend toward increased specialization of threat groups within a broader cybercrime ecosystem. In 2025, 9% of the incidents Mandiant investigated followed a "division of labor" model for the initial intrusion, up from 4% in 2022. This model involves two or more threat groups, where one group focuses exclusively on initial access operations to gain a foothold in an environment, and the subsequent group or groups perform the remaining stages of the attack lifecycle. Initial access partners, the threat clusters specializing in initial access operations, have employed low-impact, opportunistic techniques, such as malicious advertisements or non-targeted drive-by downloads to breach an organization and establish persistence. Subsequently, the initial access partner performs a "hand-off" and provides access to a secondary threat group. These secondary threat groups often engage in high-impact activity, such as deploying ransomware, during the latter stages of the targeted attack lifecycle.

Mandiant has observed evidence indicative of closer collaboration between initial access partners and secondary groups. A key signal is the time delta between the earliest activity by an initial access partner and the hand-off to a secondary threat cluster, which has been steadily decreasing since 2023. In this context, the "earliest activity" is the earliest activity in the environment that can be attributed to an attacker. This includes non-interactive events, such as the distribution of malware, and represents the moment at which a group gains access to an environment, whether they utilize it or not. A hand-off occurs at the moment of earliest activity by a secondary group. In 2022, Mandiant identified the median time between initial access and subsequent hand-off was greater than 8 hours. However, in 2025, the median time between initial access and the time at which a second group had access to the environment was 22 *seconds*. In many cases, this reflects the often automated process through which initial access partners deliver malware directly on behalf of the secondary group instead of advertising access in an underground forum.

This is when the clock starts for defenders; security teams will ideally remediate the intrusion before the moment of earliest interactive activity. The "earliest interactive activity" refers to the earliest event that signifies interactive operations, such as conducting reconnaissance or moving laterally. In some cases, the earliest activity attributed to a secondary group coincided with their earliest interactive activity. However, Mandiant has also regularly observed incidents in which a longer time delta existed between a secondary group's earliest activity and their earliest interactive activity. In these cases, activity is attributed to a secondary group before that group ever interacts with an environment, indicating that the initial access partner established access directly on behalf of the secondary group for them to utilize on their own timeline. While this still provides defenders with an opportunity to stop the attack before it becomes more consequential, the timeline is likely more constrained than if the threat actor had to sell access via underground channels.

This changing cybercrime ecosystem poses substantial challenges to security teams. Under a closer partnership model, organizations now have significantly less time to remediate low-impact initial intrusion and persistence events before a secondary group leverages that access for higher-impact activity. Organizations hunting solely for high-impact tactics, techniques, and procedures (TTPs) in their environment may not have high-priority detections in place for low-impact initial intrusion vectors. While triaging low-impact alerts in a high-volume environment places downward pressure on security teams, stopping an intrusion at a single-system stage is significantly easier than recovering from follow-on activity.

Initial Access Partners

Initial intrusion to an environment by an initial access partner is generally opportunistic and leverages low-impact TTPs. Over the last year, Mandiant has observed initial access partners utilizing Phishing (T1566), Drive-by Compromise (T1189), and Malvertising (T1583.008) techniques to later harvest credentials with infostealers, create persistence mechanisms, and conduct reconnaissance in infected environments. In some cases, an initial access partner gains a foothold into an environment prior to selling that access. This type of partner casts a wide net to build an inventory of compromised networks which they can later monetize within the broader cybercrime ecosystem. They then advertise access based on multiple factors, including the revenue of the target, industry vertical, country, and type of access being sold. Some of these groups will simply sell credentials, while others sell access to an existing malware infection or bot. Another category of initial access partner sells their ability to distribute malware or gain access on behalf of secondary groups prior to compromising an environment. These initial access partners either advertise their availability to distribute malware on behalf of secondary groups, or respond to advertisements posted by groups seeking initial access or malware distribution.

One example of an initial access partner that creates their own foothold is the financially motivated threat cluster UNC5518. In a campaign tracked by the Google Threat Intelligence Group (GTIG) in 2024 and 2025, Mandiant found that UNC5518 typically compromised legitimate websites to distribute the JavaScript downloader FAKETREFF to unsuspecting users. This campaign has evolved over time. Initially, UNC5518 injected a fake browser update lure into vulnerable websites that redirected users to threat-actor-controlled infrastructure and dropped FAKETREFF. More recently in 2025, UNC5518 began using the ClickFix technique, where a malicious pop-up instructs the user to execute the first-stage FAKETREFF payload via the Windows Run dialog or the PowerShell console. This technique bypasses the need to download any artifacts onto the user's machine directly from the browser, limiting opportunities for detection by endpoint tools and reducing the forensic footprint of the suspicious activity. Following a successful compromise, the FAKETREFF malware initiates the download of a secondary payload to establish persistence. Mandiant has identified evidence to indicate that financially motivated threat clusters such as UNC5774 purchased access to distribution services created by UNC5518 to distribute their own tooling. These secondary groups sometimes also partner with ransomware-as-a-service (RaaS) operators to obtain both the ransomware they deploy and support for the extortion of compromised organizations.

UNC6016

UNC6016 is a malware distribution threat cluster which has been active since at least November 2024. The group uses search engine optimization techniques and malicious advertising, or "malvertising", to direct users to compromised websites. Once on these sites, users are lured to download malware disguised as legitimate popular software. While this technique is neither novel nor sophisticated, it continues to be an effective means of initial access for opportunistic infections. In one investigation involving UNC6016, Mandiant identified evidence that a user was directed to a compromised website through malicious advertising and downloaded a backdoored version of a popular SSH client. Once the user executed the trojanized application, infecting their local endpoint, UNC6016 provided UNC4696 access to the resulting backdoor. UNC4696 is a financially motivated threat actor that monetizes access by deploying ransomware. Mandiant's investigation identified that this secondary impact group moved laterally, installed command and control (C2) backdoor implants, and eventually stole a significant volume of customer data. Ultimately, UNC4696 deployed ransomware across the environment and made extortion demands.

While Mandiant was investigating and helping the customer regain access to their environment, UNC6016 regained access to the same environment when an unrelated trojanized application was downloaded and executed. UNC6016 provided this foothold to another secondary group, UNC6286. UNC6286 is a financially motivated actor that almost exclusively uses UNC6016's malvertising to distribute trojanized software. In this case, UNC6286 deployed additional malware, including a dropper, a downloader, and a tunneler. The threat actors moved laterally, created new accounts, but were quickly intercepted when they attempted to acquire credentials and domain information. Unlike the initial infection where UNC4696 was able to achieve their final objective, UNC6286's activity was remediated within two hours of accessing the environment.

The swift reaction to the second intrusion was driven by the heightened vigilance of the security team during the first incident. This allowed the security team to contain the activity while the impact to the business was still minimal. However, it was the actions of the secondary group which spurred the timely response, when early detection and removal of the trojanized download could have limited impact to a single system. These intrusions highlight the necessity of a shift toward a more aggressive security posture beyond the period of active incidents, as even routine malware infections require deliberate attention and remediation to respond to security risks effectively.

activity. Additionally, investigators can expect a longer period of inactivity between the initial intrusion and follow-on activity while access is advertised and purchased.

This time delta between initial access and the earliest interactive activity by a secondary group is critical for defenders; breaking the persistent access and changing impacted passwords can prevent intrusions from escalating. In this hand-off model, the delta is, in part, determined by internal and external market forces of the cybercrime ecosystem. One such factor is the type of access being sold. For example, certain follow-on

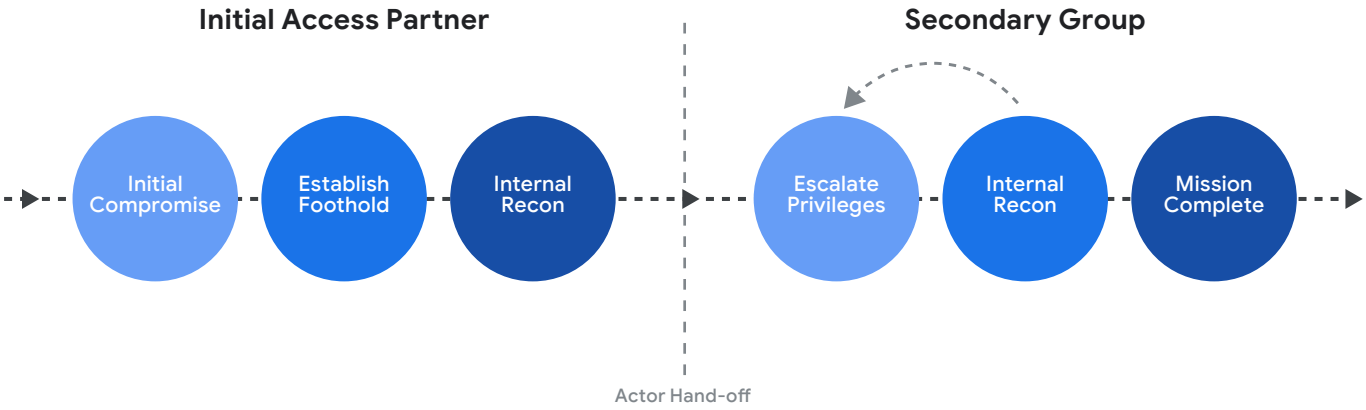

Figure 1: Initial Access Partner vs Secondary Group TTP

Handing Off

Initial access partners can follow any one of several models for handing off access, each of which manifests differently during an investigation. In the classic model, the initial access partner opportunistically acquires access to multiple environments. The group then advertises that access on underground forums with general information about the compromised organization and environment, including location, industry sector, company revenue, and type of access. Second-stage threat groups monitor those channels and purchase access in the form of IP addresses, credentials, backdoors, or other mechanisms. In these cases, there is likely to be very little overlap in TTPs between the initial access operations and follow-on

threat groups may prefer SSH credentials over installed backdoors. Some follow-on actors will refuse to purchase access to organizations in certain industry verticals, possibly due to perceived risk levels. Major cybersecurity incidents in an industry can also act as a signal to the cybercrime ecosystem to focus efforts on that industry. Additionally, high revenue in an industry may indicate to threat groups a higher reward for the potential risk. Recent changes in tradecraft, including but not limited to the availability of 0-days or the release of publicly available tooling, can also affect market conditions. Initial access partners may also internally assess operational risks such as the threat of sanctions or investigations by governments when choosing to partner with follow-on groups.

Another model for a hand-off is led by the second-stage threat group, rather than by the initial access partner. In these scenarios, a second-stage threat actor advertises on underground channels that they are seeking initial access to organizations in specific countries or possibly with specific target revenues. Some threat groups may indicate that they are unwilling to buy access to certain industry verticals, such as healthcare, government, or critical infrastructure due to operational risks. The second-stage threat groups offer a cut of the profits from a compromise, typically 20-50%, as payment for initial access. Some second-stage threat groups also advertise that they are seeking partners to distribute their own or preferred malware, which they can leverage for high-impact activity. In these scenarios, because the initial access partner does not need to find a buyer, the time between initial intrusion and earliest interactive follow-on activity is likely to be shortened. The substantial decrease in the time delta observed since 2022 between initial access and the earliest event attributed to a secondary group may indicate a rise in the prevalence of this hand-off model. Depending on the requirements of the buyer, TTPs between initial access partner and follow-on groups can either be distinct or closely linked, depending on the preferences of the secondary group.

A third, but less well defined, model for a hand-off involves initial access partners and second-stage threat groups who choose to partner with each other. Often, the exact agreements between the two or more threat groups involved are unknown, and the delineation between the groups can be blurred. In these cases, Mandiant generally observes the time between initial intrusion activity and the earliest activity attributed to a second-stage group is less than 30 seconds, consistent with the behavior of a distribution cluster. The TTPs between the two groups can often be distinct, and a close partnership can give network defenders an advantage. If a detection is raised for TTPs consistent with an initial access partner which has a known relationship to a second-stage group, the intelligence related to both groups can guide the hunt for follow-on activity.

UNC1543, UNC2165, and RansomHub

One such pair of closely linked threat groups is the initial access partner UNC1543 and the ransomware operator group UNC2165. UNC2165 is a financially motivated threat group that shares significant overlap with the threat group publicly reported as "Evil Corp," which was sanctioned by the Office of Foreign Assets Control in 2019. Although UNC1543 has provided initial access to multiple groups, UNC2165 almost exclusively obtains access to environments from UNC1543. UNC1543 typically gains access to an environment opportunistically via drive-by download distributing FAKEUPDATES, a JavaScript downloader that communicates over HTTP and is often used to install secondary payloads.

In one case investigated by Mandiant in 2025, the earliest evidence of attacker activity was traced to a browser session accessing a compromised website that led to a FAKEUPDATES infection. Mandiant observed a period of inactivity of approximately 70 minutes between the creation time of FAKEUPDATES on the target system and the hand-off to UNC2165. UNC1543 used FAKEUPDATES to drop VIPERTUNNEL, a Python-based network tunneler associated with UNC2165. The earliest interactive activity attributed to UNC2165 occurred approximately 45 minutes later, when UNC2165 used VIPERTUNNEL and reverse SSH connections to C2 addresses to establish persistent tunnels and conduct reconnaissance. The threat actors dumped credentials and moved laterally throughout the target network, then staged and stole terabytes of data through standard cloud-based file sync utilities. To complete the mission, UNC2165 destroyed system backups, then deployed RansomHub, a RaaS offering that abruptly ceased operations in April 2025, on both Windows and virtual management servers.

The use of a RaaS further demonstrates specialization within the cybercrime ecosystem. Since February 2024, multiple threat clusters, including UNC2165 and the high-impact ransomware group UNC3944, have deployed the RansomHub encryptor during incidents. This is consistent with RansomHub advertising for affiliates on underground sites. In RansomHub's business model, affiliates, who were asked to pay a $5,000 deposit, kept 90% of any ransom payments made, and the remaining 10% went to the service operators. UNC2165 historically deployed exclusive ransomware variants but has shifted to RaaS in recent years; Mandiant previously reported that this shift may have been made to impede attribution

and avoid sanctions. Partnering with specialized service providers not only reduces the technical requirements on a single threat group, but also supports anonymity.

FAKEUPDATES itself is a low-impact malware sample. The distribution method, drive-by downloads, is opportunistic and not indicative of a targeted attack. However, due to relationships such as those observed between UNC1543 and UNC2165, security teams can associate alerts for FAKEUPDATES with high criticality response efforts due to the likelihood of high-impact follow-on activity.

Recommendations

The risks posed by threat groups acting in concert to hand off access to compromised environments force a significant workload onto both security and IT staff. Many of the initial access partners rely on opportunistic infection as opposed to targeted compromise. This can result in the activity associated with initial access partners never rising to the level of concern that matches the potential impact of the secondary groups to which they provide access. The disparity between the criticality of low-impact alerts and potential high-impact outcomes necessitates a restructuring of how IT teams operate and the response playbooks used by the security teams tasked with protecting the organization while also enabling business continuity.

As the potential risk posed by low-impact activity rises, the leeway for security teams to deprioritize the response to alerts for that activity drops significantly. However, the volume of low-impact alerts and the personnel needed to review and resolve them presents a natural barrier to the effective security of an organization. Security teams are regularly overwhelmed by alert volume, leading to a condition commonly referred to as "alert fatigue" in which analysts are desensitized to alerts or alerts simply pile up until they are bulk resolved without proper analysis. This outcome can be attacked from multiple angles by creating a symbiotic relationship between security teams, IT teams, and the users that they support. Ensuring that all three groups are working towards the common goal of a secure environment that enables their work can help reduce the volume of low-impact alerts which security teams have to triage.

IT teams can integrate with security teams to deliver detailed requirements of their common workflows and help define the baseline activity which can be expected to occur within their environments. By generating a set of pre-approved and centrally stored tools for IT teams and users, security teams can help reduce the variety of process execution expected for each business unit. Either proactively deploying the tools users need or providing a low-friction interface for installation can help reduce the occurrence of opportunistic infections originating from self-installs of potentially trojanized tools. Security teams can, in turn, build detections based on the requirements provided to them by each business unit. In an environment where the pre-approved tools are provided to users, detections relating to the installation of tools or execution of binaries which fall outside of the pre-approved tooling should drop precipitously. Such detections can either be associated with potential attacker activity or provide an opportunity to better educate users as to the use of pre-approved tools defined by the requirements they provided.

Reducing the variety of the expected behavior in an environment allows security teams the necessary space to review detections of low-impact activity and respond in the time frame necessary to protect their environments. While the creation and maintenance of the baselines needed to support this outcome are time intensive, the ability to respond rapidly to risks such as initial access partners is a key factor of success for any security team. Organizations can further limit the potential impacts of a hand-off by optimizing their workflows for detection and focusing on event correlation during alert review and response procedures. Enriching alerts with contextual data can help highlight patterns of behavior which may indicate a potential high-impact intrusion pattern. Organizational alignment towards these outcomes is a crucial aspect of a mature security posture. By appropriately staffing security teams, encouraging collaboration, and increasing the signal-to-noise ratio in detection tooling, organizations create a strong foundation on which they can build. This foundation helps ensure that as the organization grows, the security infrastructure and culture necessary to achieve the desired outcomes keeps pace.

Ransomware is Now a Resilience Problem

The traditional framing of ransomware as simply a dual threat of encryption and data theft no longer captures the reality of modern extortion operations. Ransomware operators and affiliates have increasingly prioritized denying targeted organizations the ability to recover. Threat actors are targeting system and administrative planes also known as trusted service infrastructure . The terminology of "trusted service infrastructure" (TSI) is typically associated with management interfaces for platforms and technologies that provide core services for an organization such as backup technologies and virtualization platforms. This allows ransomware operators and affiliates to reduce an organization's ability to recover while maximizing the pressure to pay. They do this by attacking identity services, virtualization management planes, and backup infrastructure.

This evolution in tactics has compressed the intrusion timelines observed by Mandiant during investigations. Ransomware operators have collapsed the year-over-year dwell times for ransomware compromises, frequently seizing administrative control within hours of initial access. The velocity of these intrusions often forces defenders to preemptively disconnect their own Identity Providers (IdP) to halt the spread, effectively triggering a self-inflicted business outage to prevent total compromise. As a result, organizational survivability predicated on Endpoint Detection and Response (EDR) or traditional backup restoration at the endpoint layer are no longer sufficient recovery models. Instead, a model focused on resilience, which seeks to address the primary objectives attackers pursue, represents the best chance for organizations to keep pace with the rapid evolution of ransomware operators.

In this context, the TSI fabric is not simply an inventory of "critical assets." It is the operational terrain on which ransomware campaigns achieve scale and where recovery is either enabled or prevented.

Attackers Weaponize the Administrative Fabric

Modern ransomware campaigns are no longer random smash-and-grab operations; they are, instead, structured takeovers of the IT administrative infrastructure. Attackers understand that an enterprise is an interdependent web of trust. In 2025, Mandiant investigated multiple breaches in which threat actors used "Living Off the Land" techniques, turning an organization's own administrative tooling and trusted security controls into primary attack vectors.

The Identity Core

The path to total control of a targeted environment almost always begins with identity, but sophisticated groups have moved beyond simple credential theft and are, instead, manipulating the identity control plane. In many of the ransomware incidents which Mandiant investigated in 2025, attackers exploited misconfigured Active Directory (AD) Certificate Services templates to issue certificates and create or impersonate admin accounts. These attacker-controlled accounts weren't subject to multi-factor authentication (MFA) and were excluded from password rotation requirements. Mandiant has also investigated incidents where attackers used administrative commands and specialized extraction tools to steal the entire AD database, exposing every hash and effectively compromising every user credential in the domain.

Threat actors have also started to engage in forms of high-volume credential harvesting and disruption activity during intrusions. In several investigations, Mandiant identified evidence that the threat actors targeted the organization's domain controllers (DC) and then accessed localized enterprise credential vaults to extract dozens of high-privilege credentials in a single session. To further destabilize the environment, adversaries executed rapid reconnaissance followed by scheduled tasks which forced password changes for privileged administrative accounts. This combination of actions locked defenders out of their own emergency accounts during a crisis.

In some attacks, the targeting of identity extended to the cloud. Adversaries pivoted from a compromised Microsoft Entra ID tenant to hijack associated Azure storage accounts, utilizing the corporate identity provider to unlock and destroy the infrastructure backend.

Fleet Management

Many of the tools on which organizations rely for day-to-day operations often become targets during ransomware events. Ransomware threat actors have extended this targeting to even include the tools which were designed for patching and security due to their ability to distribute packages at scale within the environment. Mandiant has identified evidence of threat actors weaponizing group policy objects (GPO) to create scheduled tasks across thousands of endpoints which execute immediately, turning the domain's own management fabric into a mass-distribution engine for ransomware. In other instances, attackers compromised Microsoft Endpoint Manager service accounts, leveraging their administrative privileges to pivot directly to a DC and seize control of the entire domain hierarchy. In some cases, threat actors use the organization's fleet management capabilities to disable features, such as PowerShell logging, which can be critical to attacker detection across the environment.

The Virtualization Layer

In recent investigations, Mandiant has identified attacker behaviors that indicate threat actors are treating the hypervisor as more than a target for encryption. Rather than only encrypting the hypervisor and guest operating systems, threat actors are also using it as a staging ground for data theft. In recent investigations, Mandiant identified evidence of threat actors compromising virtualization hosts to execute reconnaissance tools and archiving utilities directly on the virtualization host. Attackers used archiving tools directly on the hypervisor to compress and archive the Virtual Hard Disks of backup servers. This allowed the attackers to steal critical data volumes while bypassing data loss prevention controls running inside the guest operating systems.

Adversaries also exploit the convergence of Identity and Virtualization to seize control. Mandiant has identified ransomware threat actors weaponizing the default virtualization admin group used in some AD integrations. By adding a compromised service account to the AD group, they automatically inherited administrator privileges on all domain-joined virtualization hosts. This allowed the threat actor to disable firewalls and mass-deploy ransomware to dozens of hypervisors instantly. The impact of targeting the virtualization layer can often be comprehensive. An Akira ransomware attack resulted in the successful encryption of most virtualized servers. This resulted in a complete loss of investigative visibility due to network telemetry logs having been encrypted alongside the production servers. The loss of critical telemetry created substantial friction in the organization's efforts to identify data theft.

Targeting Safety Nets

The final and often most devastating stage of the compromise is the destruction of backups. Attackers actively perform reconnaissance on backup architectures, accessing admin consoles and documentation repositories to map out storage locations, SQL configurations, and retrieve encryption configurations. Once the location and access to resources are mapped out, threat actors will often systematically destroy the elements critical to the recovery of the environment. The information gleaned from the reconnaissance phase of the intrusion can also form the basis for the escalation of privileges necessary to impact backup infrastructure. In one example, Mandiant identified evidence that the threat actor compromised backup management servers in order to extract and decrypt credentials which were stored in the configuration database. The exposed credentials granted the threat actors direct access to the administrative passwords for storage controllers and cloud accounts. The threat actors used the stolen credentials and keys to execute broad-scale deletion commands, wipe millions of backup objects from cloud storage, and delete dozens of local system backups.

In on-premises scenarios, the deletion activity can be equally as destructive and thorough. Mandiant identified evidence that the attackers deleted the backup configurations from the backup server which unlinked the virtualization environment from the backup platform. While the backup data was held on immutable storage and not destroyed, it became inaccessible through the backup platform user interface and required a lengthy engagement with the vendor of the backup solution to start recovery. In another example, threat actors compromised network-attached storage (NAS) appliances via SSH to enable the local admin account and change passwords. The attackers then installed data theft tools to steal critical volumes before encrypting those same volumes, ensuring that the local recovery point was destroyed.

Where Recovery Can Fail

The compromise and weaponization of Tier-0 assets transcend a simple service outage; it creates a fundamental recovery deadlock during which the foundational trust required for restoration is lost. In such scenarios, the standard incident response playbooks break down because the tools and trust required to execute them are unavailable. Mandiant's frontline experience highlights three primary failure scenarios that can rapidly escalate the criticality of a ransomware incident.

Identity

When the identity fabric is fundamentally compromised, such as via Windows NT directory service theft or AD Certificate Service forgery, the organization faces a systematic identity collapse. Defenders cannot simply reset passwords because the attacker possesses the cryptographic keys necessary to complete the reset. This can often force a "Greenfield" recovery during which a new, trusted AD forest is created while operations remain impacted in the production environment. This form of recovery imposes a staggering operational tax as restoration timelines stretch from days to weeks, and financial losses compound as IT teams are forced to rebuild the entire identity backbone manually rather than simply restoring data. Crucially, this often extends to organizational communication channels as well. If the compromised identity provider federates to corporate email or collaboration platforms, defenders lose their primary means of coordination during response activity. As a result, they are forced to rely on out-of-band communications or risk leaking sensitive information to an attacker who retains broad access to the collaboration and email platforms.

Velocity

The encryption or destruction of virtualization and backup control planes can severely impede the progress of restoration activity in impacted environments. When an organization's hypervisors are encrypted alongside the on-premises backups, the hardware may still function, but the logical infrastructure is inoperable. In cases where backup catalogs are wiped or cloud buckets are deleted via stolen keys, the organization loses its primary recovery capability. As a result, teams working to restore the environment are forced to rely on slower, older archival data, leading to a risk of additional data loss and extended downtime.

Re-Compromise

The recovery of compromised environments often relies on workflows where teams reimage impacted systems and restore data from backups. These workflows, in turn, rely on the availability of trusted restoration sources. A critical aspect of recovery denial activity is the embedding of persistence mechanisms into the infrastructure used to manage the recovery. By configuring fleet management tools or hypervisor templates to include the deployment of attacker-controlled backdoors as part of deployment operations, threat actors ensure they are able to retain access to the targeted environment. Consequently, the very process of remediation triggers a reinfection of the environment, turning the organization's own recovery tools into a distribution mechanism for the threat actor.

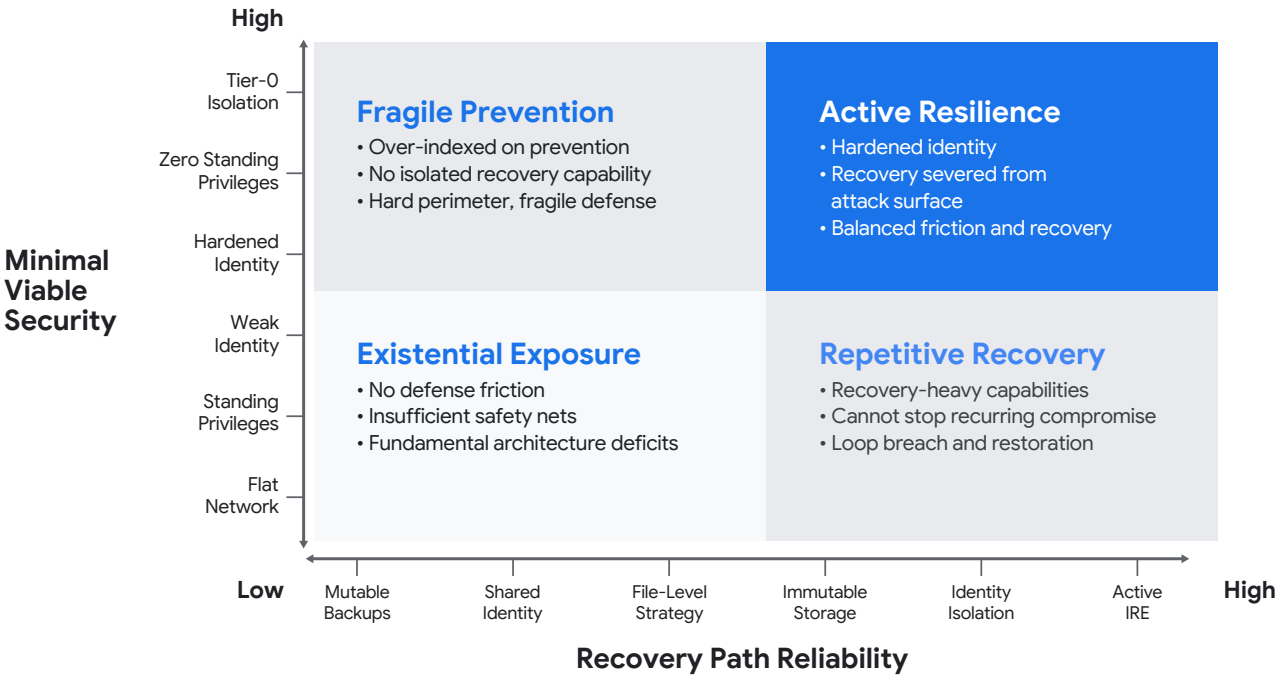

Figure 2: The Resilience Maturity Matrix

Implementing Active Resilience

The Resilience Maturity Matrix

Resilience is not a single metric; it is the product of **prevention** and **recovery**. Prevention can be measured as the degree of friction a threat actor experiences in trying to move through a targeted environment. Recovery can be measured by the level of confidence an organization has in the restored environment. The controls which influence these values can be mapped against two primary axes; Minimal Viable Security and Recovery Path Reliability.

The first quadrant, Existential Exposure, represents a state where an organization operates with fundamental architectural deficits. A lack of defensive friction allows adversaries to move unimpeded within the environment, while the absence of offline safety nets ensures that the recovery environment is likely destroyed alongside production. The outcome is often severe; recovery is statistically unlikely without acceding to extortion demands, and even then, success is not guaranteed.

In contrast, Fragile Prevention occurs when organizations are over-indexed on prevention but lack ancillary recovery infrastructure such as isolated recovery capability, immutable storage, or off-network copies of data. While the perimeter is difficult to breach, the defenses are rigid and fragile. Consequently, a single control plane bypass can result in an unrecoverable event because the recovery fabric shares the fate of the production environment.

In Repetitive Recovery, the organization has invested heavily in safety nets but lacks the defensive friction required to contain attackers after the initial breach. While they can restore their data, they cannot stop the recurring compromise. This leaves the organization able to survive the initial event but suffering high operational costs and continuous downtime cycles; trapped in a loop of breach and restoration.

Active Resilience is the target state, characterized by balanced friction and trusted recovery. The adversary is slowed by hardened identity controls, and the recovery fabric is architecturally severed from the attack surface. In this state, attacks are degraded from existential crises into containable incidents. The organization retains the authority to refuse extortion demands, secure in the knowledge that their recovery path remains untainted.

Active Resilience

To move from Existential Exposure to Active Resilience, organizations should aim to increase defensive friction and isolate the recovery path.

Increase Minimal Viable Security (MVS)

Focus: Hardening Identity, Network, and Virtualization Layers.

Secure Identity and Device Perimeters

Identity is the fastest path to TSI control, which is why modern ransomware operations aim for control of privilege mechanisms and not just passwords. Conditional access and Privileged Identity Management (PIM) matter because they make administrative access time-bound, justified, and observable rather than ambient. Hardening high-impact identity pathways, such as auditing SPN accounts for Kerberoasting exposure and maintaining a practiced process to reset the KRBTGT twice during incident response, helps increase the friction of attacker movement and confidence in restoration. Administrative endpoints are part of the control plane. Using hardened Privileged Access Workstations (PAWs), enforcing MFA and strong authentication, and reducing lateral movement shortcuts increases the attacker's cost to move from foothold to domain-level control.

Network and System Hardening

Ransomware resilience improves when environments are intentionally difficult to traverse quickly. Zero-Trust segmentation that isolates TSI limits lateral movement into the administrative fabric and slows the sequence that typically precedes large-scale encryption or backup sabotage. Decommissioning end-of-life systems, protocols, and configurations can remove common footholds and reduce blind spots. Telemetry should be tuned to the "Living Off the Land" behaviors, particularly partial WMI execution and unauthorized PowerShell activity, so defenders can identify the administrative takeover phase before it impacts the environment more broadly.

Hypervisor Hardening and Tier-0 Isolation

Virtualization control planes are force multipliers for threat actors; a single compromise can translate into broad operational failure. Formally treating hypervisors and their management planes as TSI ensures the strictest access constraints and highest-priority monitoring. Severing AD integration reduces the risk that identity compromise becomes instant hypervisor-wide control, and dedicated out-of-band management with local, MFA-protected accounts reinforces that separation.

Advance Recovery Path Reliability (RPR)

Focus: Survivability and Identity Isolation.

Survivable Recovery and Disaster Planning

Recovery often fails when trust anchors are compromised alongside production, especially identity. Offline or immutable versions of Tier-0 assets create a recovery option that doesn't share the fate of production when attackers delete backups, corrupt catalogs, or seize admin consoles.

A dedicated, isolated recovery environment allows teams to clean, validate, and stage system restores before reconnecting to production. An isolated recovery environment can directly mitigate a reinfection loop commonly found in the Second Wave failure mode. This requires a formal Disaster Recovery process that assumes loss of the primary site and primary identity fabric, to account for ransomware operators increasingly designing their operations to target those resources as a means of denying recovery.

Tactical Governance and Threat Intelligence

Static policies can become liabilities when attackers adapt to regain access or sabotage restoration. Updating verification procedures based on emerging tactics, techniques, and procedures, such as tightening help-desk workflows when social engineering patterns shift, helps reduce reentry opportunities during the most fragile phase of recovery. Governance is of equal importance during a recovery. Organizational leaders must proactively define decision recovery playbooks that prioritize integrity of the recovery fabric over short-term uptime. This enables responders to preserve trusted recovery even under pressure.

Resilience as a Contested Domain

The emergence of recovery denial marks a strategic inflection point in the landscape of cyber extortion. Historically, organizations architected defenses under the assumption that, while production environments were contested, recovery environments remained secure. Adversaries have industrialized the targeting of Tier-0 administrative infrastructure, recognizing that compromising an organization's ability to recover is as effective a lever for extortion as the theft of data.

In this hostile landscape, the convergence of Information Security and Business Continuity is absolute. A strategy reliant solely on prevention is inherently fragile as once an adversary renders the infrastructure unresponsive to control, traditional detection stacks offer little in the way of mitigation. Resilience requires an architecture designed to endure and recover from the compromise of these internal control planes.

Ultimately, the transition to Active Resilience is achieved through architectural discipline rather than tool acquisition. It necessitates the strict decoupling of critical systems, the implementation of friction to impede lateral movement, and recovery abilities that operate independently of the primary identity fabric. Shifting towards Active Resilience enables leaders to do more than just secure their data; it helps ensure an organization can maintain their operations while under adversarial pressure.

Multi-Year Intrusions Highlighting Extreme Persistence

Sophisticated threat actors seeking long-term access to targeted environments remain a persistent and effective threat. As threat actors continuously evolve to circumvent modern security technologies, the risk of covert attacker activity going unnoticed increases. While these adversaries may deploy custom malware when necessary, they frequently prioritize stolen credentials, limit their tooling to what is available in the targeted environment, and target systems less likely to be instrumented with security tooling. Taken as a whole, state-sponsored espionage threat actors represent heightened risk while presenting fewer and fewer opportunities for detection. Recent activity, such as UNC6201's BRICKSTORM campaigns, and UNC5807's targeting of telecommunications networks, illustrates the ongoing adaptability of sophisticated threat actors. Similarly, UNC1549's targeting of third-party accounts to compromise their intended targets demonstrates the concerted efforts espionage threat actors may engage in as they seek long term access. Mandiant has also identified evidence of suspected state sponsored threat groups using anti-forensics techniques to further obscure the activities they take in an environment. In all four campaigns, these suspected espionage threat actors masqueraded as legitimate users within the target environment, and consistently pursued efforts that focused on long-term access. As a result, in all four campaigns, the threat actors were able to maintain multi-year intrusions into some of the targeted environments. Intrusions that span extended time periods force defenders and investigators to rely on data sources that are historically under-utilized to answer critical questions regarding the security of the environment and the actions of the attacker.

Observations

The operational discipline required to sustain multi-year intrusions is perhaps best exemplified by UNC6201's BRICKSTORM campaign. UNC6201 is a suspected PRC-nexus espionage threat group that uses a sophisticated backdoor called BRICKSTORM to target virtualization infrastructure. While notable overlaps in UNC6201 and UNC5221 activity have been reported as synonymous with the threat actor "Silk Typhoon" in public reporting, Google Threat Intelligence Group (GTIG) does not currently consider the two clusters of activity to be the same. With an average dwell time exceeding 390 days, UNC6201 established persistent footholds in compromised environments using BRICKSTORM, and often focused follow-on activity towards network edge appliances that sit outside the reach of endpoint detection and response (EDR) tools. This threat actor in particular focuses on gathering credentials by cloning sensitive virtual machines, such as Domain Controllers and Privileged Access Management (PAM) systems. By securing a cache of valid credentials, UNC6201 was able to blend in with legitimate activity, effectively hiding their malicious actions as potential system administrator and service account activity. The combination of legitimate credentials and operational discipline that minimized detection opportunities enabled UNC6201 to ensure pervasive access to targeted environments for extended periods of time.

Similar to UNC6201, the PRC-nexus cyber espionage group UNC5807, likely active since 2020, tends to focus its efforts on telecommunication backbones and network routing infrastructure. UNC5807's activity is generally consistent with the threat actor discussed in public reporting as "Salt Typhoon." UNC5807 systematically exploits vulnerabilities in exposed network edge devices such as VPN solutions and routers to establish initial access. Upon establishing an initial foothold, UNC5807 commonly targets authentication protocols such as TACACS+ and RADIUS for collection using packet capture functionalities native to the appliances they compromise. To sustain these complex operations, UNC5807 uses a specialized, purpose-built toolkit designed to bypass the limitations of traditional security telemetry on non-standard appliance platforms.

While some actors focus on the hardware edge, a suspected Iranian cluster of cyber espionage activity, tracked by GTIG as UNC1549, exploits the boundaries of organizational trust. Primarily targeting the global aerospace and defense sectors, UNC1549 uses highly targeted spear phishing lures to gain a foothold into environments. They frequently pivot through less well-defended third-party service providers, taking advantage of established network connectivity to reach their intended targets. In one intrusion Mandiant investigated, UNC1549 authenticated to an instance within a customer's Virtual Desktop Infrastructure (VDI) using stolen credentials. The desktop instance UNC1549 targeted was originally provisioned to provide access to a third-party vendor. Once authenticated, the threat actors bypassed process launch policies configured to limit the types of applications that could run on the instance. By bypassing the controls on the instance, UNC1549 was able to pivot into the environment of their intended target.

Lastly, Mandiant identified a suspected state-sponsored threat actor using anti-forensics and in-memory persistence techniques to reduce the opportunities for detection related to their activities. During recent investigations, the threat group demonstrated effective anti-forensics techniques by using built-in tools to remove evidence of their activity from system resources such as the Linux WTMP login records and entries in the Audit system log files. The threat group also used timestomping techniques to mask file system modification timestamps of the files with which they interacted or altered. To maintain a minimal on-disk footprint, the threat actor frequently deleted malware binaries after execution, allowing the process to run exclusively in volatile memory on systems that underwent infrequent reboot cycles. Limited EDR capabilities and restricted network logging also contributed to the threat group's ability to remain undetected for multiple years.

Challenges

Investigations into threat actor activity with long dwell times often highlight several challenges to early detection and efficient incident scoping. Gaps in network visibility and telemetry, asset management inefficiencies, and limitations in data retention can all contribute to the operational friction experienced by organizations as they seek to identify and remediate a compromise of their environment.

Telemetry Gaps

While many organizations have invested in better visibility into their environment through EDR telemetry and network logging, few have opted to retain that data for long-term analysis. In many environments, logs that are either native on hosts or forwarded to SIEMs do not extend far enough back in time to scope initial access if a multi-year compromise is identified in their environment. As an example, the average dwell time of a case in which BRICKSTORM was deployed was 393 days, but many organizations only log data for 90 days due to cost-saving measures. In some cases, organizations opt to keep logs for 360 days but these would still be insufficient to scope the full activity of the average BRICKSTORM campaign. In such instances, evidence likely will not exist to determine the initial access vector, which severely impacts the confidence an organization can have in remediation activity.

EDR tooling acts as an enabler for security teams in the ongoing defense of their environments, but EDR solutions present similar challenges with regards to deployment and retention. While EDR alerts are often retained for 180 days or more, real-time telemetry is typically limited to the 7 to 30 day range. Long-term retention of telemetry data within the EDR platform can be prohibitively expensive and impact the responsiveness of the analysis platform in ongoing monitoring. Filtered real-time telemetry produced by EDR solutions that is copied to long-term storage can provide a secondary high-fidelity source of data if a compromise that extends past the retention window is identified. By partnering with infrastructure and operations groups, internal security teams can define a set of event types that correspond to the most critical questions raised during an investigation. Events that detail authentication behavior, process execution, internal access to sensitive resources, and privilege escalation are a good starting point for collection. Including events such as registry modification, service and scheduled task creation, and command script details helps add additional context and depth to the set of events forwarded to long-term storage. Security teams can then assess the volume of data being collected as well as the frequency of individual event types to identify potentially noisy feeds. Partnering further with infrastructure and operations teams can help determine the bounds of expected events generated from those environments and create additional tuning rules for the filtered logs.

An additional obstacle is the common lack of visibility into activity within the network. While most organizations place a premium on visibility into external activity targeting the perimeter, many give the internal environment far less attention. The absence of inter-network traffic logging or, alternatively, system authentication logging beyond Windows Domain Controller Security event logs, prevents investigators from accurately mapping lateral movement between internal systems and network segments. This problem is further compounded as threat groups learn to mask malicious behavior and blend into the environments they compromise. Limited logging also hinders the ability to identify a threat actor's specific targets and objectives. In one case that Mandiant investigated, a client's logging configuration resulted in all cloud-bound traffic from a compromised on-premises segment appearing as though it originated from a single egress IP. The lack of visibility into the on-premises segment made it impossible to isolate traffic from known compromised hosts. Failing to log access to high-criticality systems or those that proxy connections on behalf of others can constitute a significant telemetry gap.

Prioritizing the forwarding of logs to a centralized location can help improve the retention of telemetry and secure them from anti-forensics activity. This also helps secure the integrity of log sources against threat actors who seek to wipe or modify log sources on the hosts they compromise. Furthermore, logging repositories should be configured to generate alerts if a source of logging is unexpectedly terminated. These alerts provide operations teams the opportunity to correct logging misconfigurations or, potentially, identify the intentional disabling of logs by a threat actor. While the duration of logs stored is a critical metric, the scope of logging sources should also be audited and, where possible, improved to cover frequently overlooked assets. Logs from hypervisors, network edge devices, and routers, as well as authentication data for all Windows and Linux servers in an environment, help reduce the areas in which long-term intrusions can hide. Finally, logging configurations should be audited periodically to ensure logs are not only being captured, but that the appropriate level of detail is being provided.

Asset Blind Spots

Asset management is a fundamental aspect of a mature cybersecurity program; however, it is frequently not implemented with the rigor needed to support complex investigations. During complex investigations, gaps in asset management are often quickly highlighted. These gaps most commonly manifest when network defenders attempt to identify an asset associated with potential malicious activity, only to find that no centralized source of truth exists. The time and resources required to track down the system in question injects significant friction into the process. Even in cases where an organization has retained sufficient log data to support the analysis, the inability to pivot quickly from a finding to the identification and analysis of the system impacts the speed at which network defenders can respond to potential threats or scope long-term compromises.

Gaps in an organization's management and tracking of their assets can occur for a variety of reasons. Appliances such as network edge devices, routers, and hypervisors are often excluded from asset management. These specialized appliances frequently lack support for conventional monitoring and security tooling, which can make them prime targets for sophisticated threat actors seeking long-term covert access to an environment. This often results in diminished visibility into these assets for network defenders, which can compound other gaps in visibility when the assets are unknowingly excluded from security and logging improvements.

Additionally, as organizations grow in size, the challenges of asset management increase if a strict culture of organizational collaboration doesn't exist. If individual business units make decisions with respect to infrastructure without the cooperation of the larger organization, they can introduce systems that lack appropriate instrumentation and asset tracking. The ability to create an accounting of unmanaged systems in an environment during an investigation into suspicious behavior, or worse, while investigating a multi-year intrusion, is a costly and time consuming endeavor. A collaborative culture that demonstrates the value of interdependency between business units and security teams can help ensure answers to critical questions are delivered in a timely manner while also helping reduce visibility gaps.

Recommendations

Much like any emergency, the worst possible time to discover that your response preparedness is lacking is during a live crisis. Maintaining the velocity of the response to a successful attack is crucial to a positive outcome and the recovery of the environment. As investigators and security teams work to scope the intrusion and ultimately remove the attacker from the environment, many of the assumptions which internal teams have of their systems and workflows are tested and it is not uncommon to find gaps in knowledge, instrumentation, and logging. Visibility gaps, security instrumentation deployments, or even simple infrastructure documentation such as network diagrams can slow an investigation dramatically as critical resources are pulled from the response to find answers that the security teams require. A compromise in which threat actor activity spans multiple years compounds these issues while also raising difficult questions. In lacking adequate data to detail the scope of an intrusion and potential data theft, organizations risk a loss of trust with customers during disclosure and can face scrutiny relating to breach disclosure laws. Without the ability to quantify data theft, organizations may find they are forced to assume the worst.

While prevention of a compromise is an ideal outcome for any targeted attack, the preparation required to respond to the successful attacks requires long-term support and planning. User education which focuses on turning users into assets that can help drive positive security outcomes and open collaboration between security teams and the business units they support provide the best foundation for building mature security programs. A collaborative effort to ensure the organization has an appropriate accounting of and visibility into the systems on which they rely also helps set the stage for active validation of the results.

Organizations should regularly test the assumptions their teams have of the environment to identify and validate the data. During these audits, efforts to ensure complete coverage of the environment should coincide with an ongoing effort to identify and remove unused technology and systems to reduce administrative burden and risk. By working to overcome the limitations of standard logging and EDR through a proactive review of total available telemetry, identifying and closing gaps, and maintaining appropriate detail in documentation, defenders can more efficiently and effectively respond to incidents. Doing this prework enables defenders to move beyond atomic IOC searches, which are less effective against threat actors that proactively avoid techniques susceptible to static detections. It allows them to focus on more advanced analysis techniques like stack ranking data to look for outliers and comparison of data against known good systems. These techniques are much more effective in identifying sophisticated threat actor activity.

Organizations should also implement a routine of proactive threat hunting. Relying solely on reactive alerts generated by conventional security tools is insufficient for detecting sophisticated, evasive threat actors. Proactive hunting serves to bridge this critical visibility gap and can significantly curtail the dwell time of an adversary within the network. Since threat actors frequently use legitimate, native administrative utilities to hide their actions in plain sight, detection hinges on the analyst's ability to distinguish benign administrative activity from an attacker's malicious actions. Furthermore, integrating the latest threat intelligence into the hunting process allows organizations to focus on specific techniques currently employed by adversaries. Ultimately, the systematic practice of proactive threat hunting not only identifies existing breaches but also highlights and enables the refinement or remediation of operational deficiencies in the defense posture.

Adversary Focus on Virtualized Infrastructure

In recent years, threat actors have increasingly targeted virtualized infrastructure to accomplish their objectives in a targeted environment. Virtualization platforms often consist of three main components: a dedicated, centralized management server to administer virtual machines, hypervisors to handle hardware allocation, and the virtual machines themselves. Mandiant has identified threat actor activity targeting each of the three components of virtualized infrastructure throughout the attacker lifecycle–from achieving initial access to stealing sensitive data.

Weaponizing the Virtualization Stack

Evasion of Security Tooling

Modern defenses for the virtualization stack often rely on Endpoint Detection and Response (EDR) solutions designed to protect the guest operating systems on which they are installed. The hypervisors themselves, however, may run specialized proprietary operating systems that are incompatible with modern EDR solutions. This creates a blind spot for security teams and a potential means through which threat actors can pursue their objectives while remaining undetected. Recent data substantiates this claim; Mandiant identified several malware families capable of targeting virtual infrastructure, including REDBIKE, commonly known as "Akira" ransomware. In 2025, REDBIKE was the malware family most frequently identified by Mandiant during ransomware investigations.

In some investigations, Mandiant identified evidence of threat actors not only targeting virtualized infrastructure but also creating new unmanaged virtual machines to support ongoing operations in the targeted environment. These virtual machines existed outside of the organization's asset management and security monitoring systems, which enabled the threat actor to conduct deeper staging operations without being flagged by EDR or SIEM alerts. The existence of unmanaged and unaccounted for virtual machines can also result in severe delays when security teams are attempting to investigate suspicious behavior. The overhead in identifying an unknown system while trying to triage potential attacker activity can impact the velocity of the investigation or, worse, lead to dead ends where high-confidence determinations cannot be made.

Privilege Escalation

A critical advantage in attacking the hypervisor, which hosts a system of operational interest to the attacker instead of the system itself, is the ability to steal sensitive data without directly interacting with the target system. This tactic can often be deployed by a threat actor as part of their privilege escalation techniques once they've gained access to the hypervisor. A common path for privilege escalation includes stealing the ntds.dit file, which is the core database file for the Active Directory (AD) Domain Services. To accomplish this, threat actors compromise a Domain Controller, elevate their privileges, and often run additional tooling in order to steal the ntds.dit file. During each of these operations, threat actors create opportunities for detection especially as the access to ntds.dit is a common indicator of attacker activity supported by all modern EDR solutions. However, in an attack where the domain controller is part of virtualized infrastructure, threat actors can simply target the virtualization storage layer directly. Using a technique commonly referred to as "Snapshot Mounting," adversaries create a snapshot of a running Domain Controller, clone or mount that snapshot's virtual disk (VMDK) to a rogue virtual machine, and extract the ntds.dit file. By targeting the virtualization storage layer, the EDR on the guest operating system remains unaware of the access to the underlying files. Similarly, the hypervisor's management of physical resources such as memory presents another avenue through which threat actors can escalate privileges. By targeting the individual virtual memory files through a compromised hypervisor, threat actors can transfer those files out of the environment, where they are parsed by tools such as Mimikatz to extract credentials.

While direct interaction with resources managed by the hypervisor represents a high level of risk, threat actors have also been found to rely on the same hypervisor management platforms and interfaces that the organization uses for day-to-day operations. In environments where the hypervisor and management server are integrated with AD and domain-joined, threat actors can exploit this trust relationship to achieve unauthorized administrative control. Specifically, by compromising a privileged AD account, an attacker can create or manipulate the AD group responsible for provisioning rights across the virtual infrastructure to grant themselves full administrative access.

Persistence

Adversaries frequently target the virtualization platform to establish persistence within an environment through access to the hypervisor's shell. By executing commands directly within this environment, they can deploy persistence mechanisms that survive standard remediations and system reboots. For instance, Mandiant identified evidence to indicate attackers exploited vulnerabilities to upload web-based backdoors, effectively turning the web management interface itself into a persistent gateway for command execution. Others established persistence by registering malicious systemd services. Systemd is the default system and service management interface for Linux distributions. Creating a malicious service ensured that the threat actor's backdoors automatically restart whenever the host reboots.

A more sophisticated example of targeting the hypervisor for persistence involved the use of an open-source machine emulator, QEMU, launched as a raw process on the hypervisor. In this example, the threat actor created an unmanaged VM and installed command and control frameworks alongside tunneling utilities to create a secure backdoor. Because this VM was running on the hypervisor instead of being run by the hypervisor, the virtual instance failed to show up on listings of guest instances. This allowed the VM to not only avoid detection through EDR solutions and scanning, but also obscure the origin point of the instance itself.

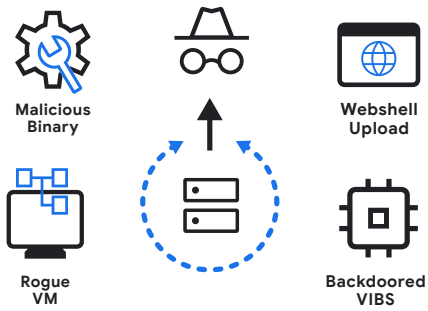

Figure 3: Ransomware impact

Ransomware Impact

To combat the impacts of ransomware incidents, organizations have developed comprehensive recovery plans and backup solutions. As such, more recent ransomware campaigns prioritize the destruction of backups to eliminate recovery options and increase leverage during extortion demands. Many enterprise backup solutions offer plug-ins to virtualization management servers and virtual machine architecture to facilitate easy restoration from backups. In two recent cases, Mandiant identified evidence that the threat actor compromised the service accounts used for the communications between the backup provider and the virtualization platform. The threat actor then used the compromised service accounts to destroy backup and restoration resources, limiting the ability of the targeted organization to recover independently.

Threat actors also prioritize stealthy ransomware deployment methods to avoid automated detection and response from EDR tooling. Modern EDRs have the capability to respond in real time to detect, quarantine, and block the execution of ransomware binaries themselves or the actions commonly employed by ransomware threat actors. The virtualization stack, however, presents threat actors with the opportunity to deploy ransomware at the hypervisor level. Rather than deploying an encryptor to individual virtual machines that are more likely to have EDR tooling installed, the threat actor can encrypt the datastore files, such as .VMDKs and .VMEMs, stored on the hypervisor. Encrypting these datastore files on the hypervisor renders any associated virtual machine inoperable and unrecoverable without decryption or an unaffected backup. In 2025, Mandiant identified evidence

of threat actors deploying and executing custom malware such as FOULFOG.LINUX or INC.LINUX on compromised hypervisors. Mandiant's analysis of recovered malware samples revealed that on execution, these binaries shut down the virtual machines running on the hypervisor and explicitly targeted virtual machine datastore files for encryption.

Logging Guidance

Virtualization platform logs can provide high-quality detailed data regarding the operations and use of the virtualization tools in an environment. Unfortunately, Mandiant has found that these logs are often excluded from enterprise logging solutions. By forwarding virtualization platform logs to a centralized location, such as a SIEM, security teams significantly enhance their ability to identify and alert on malicious activity that targets the management server and hypervisors. This can help close the gap in visibility into threat actor activity occurring outside the standard guest OS monitoring.

The Control Plane

Management server-level logs can provide a high-level narrative of attacks targeting the platform by recording identity-driven events, such as logins, task initiations, and permission changes. These platforms typically utilize a dedicated application-layer log source, distinct from the underlying operating system, to track user actions, permission changes, and authentication events.

Authentication and Authorization	VM Lifecycle and State Changes	Configuration and Inventory Changes
Logon/Logoff Events: Track high-level login successful authentication and failures.	Power Operations: Power On, Power Off, Reset, and Suspend operational logs.	Hardware Reconfiguration: Adding/removing vCPUs, changing RAM allocation, or expanding virtual disks.
Role/Permission Changes: Logs indicating that a user was granted "Administrator" privileges or added to a sensitive group.	Snapshot Activity: Creation, deletion, and consolidation of snapshots.	Network Changes: Modifications to virtual switches, port groups, or VLAN tagging.
API Access: If you use automation (Terraform, Ansible), log the authentication of service accounts accessing the API.	VM Creation/Deletion: Logs capturing the deployment of new or removal of VMs.	Storage Changes: Mounting or unmounting datastores/LUNs.
Web Request Logs: These logs trace incoming web requests, including the source IP address connecting to the web client.	Migrations: Logs showing VMs moving between hosts.	Host Operations: Putting a host into "Maintenance Mode" or adding/removing a host from a cluster.

The Data Plane

While management logs offer a high-level narrative, hypervisor-level logging captures the granular execution of the kernel and the physical state of the underlying hardware. This provides a distinct telemetry source for monitoring direct host shell access and low-level networking events to address critical visibility gaps in detection and response workflows.

Host Authentication and Access	Privilege and Integrity Monitoring	Network and Isolation Violations
SSH Remote Access: Logs capturing remote logins via SSH to the hypervisor.	Privilege Escalation: Logs recording the use of su (substitute user) or sudo commands within the host shell to elevate permissions from a standard user to root.	Promiscuous Mode: Logs indicating a VM's vNIC was set to "Promiscuous Mode" allowing it to sniff all traffic on the virtual switch.
Direct Console/Shell Access: Logs showing interaction with the physical TTY, KVM, or Direct Console User Interface (DCUI).	Unsigned Package/ Module Loading: Alerts when an administrator attempts to force-install an unsigned driver, kernel module, or VIB (vSphere Installation Bundle), which can be a vector for rootkits or malware.	MAC Address Impersonation: Logs detecting a VM attempting to send traffic using a Source MAC address different from its assigned one.
Direct Web/API Access: Logs capturing the Web/ API authentication directly to the hypervisor.	Host Configuration Tampering: Logs capturing manual edits to critical system configuration files (e.g., /etc/passwd, esx. conf) performed via the command line.	Host Firewall Modification: Changes to the hypervisor's local firewall rules (e.g., opening unexpected ports like HTTP or Telnet) to expose the management interface to the open internet.

Proactive Security Recommendations

The security of an organization's virtualization platform defines the security of every VM it supports. When hosting Tier-0 services like Active Directory or Secrets Management, the underlying virtualization environment must be hardened as a Tier-0 asset. Without strict isolation and equivalent protection levels, these critical systems are only as secure as the management plane on which they run.

Identity Modernization

Configuring hypervisors as domain-joined hosts can present inherent risks to an environment. Where possible, organizations should review the practicality of removing hypervisors from their domain to prevent threat actors from using a compromised domain to gain administrative control over a virtualization platform. Phishing-resistant Multi-Factor Authentication (MFA)

such as FIDO2 for all hypervisor logins and privileged accounts can reduce the threat of credential theft and lateral movement.

Organizations can further reduce their attack surface by deploying a dedicated Infrastructure-Only Identity Provider (IdP) or a hardened Privileged Access Management (PAM) solution. A dedicated IdP ensures that infrastructure administrative accounts exist in a separate, restricted directory that is entirely independent of the corporate domain. Alternatively, a PAM-centric approach can provide strong controls such as enforcing Just-in-Time (JIT) access. This ensures that administrative privileges are granted only on a temporary and as-needed basis, which helps reduce the availability of persistent high-level access threat actors seek to exploit for long-term presence in the environment.

Stack Hardening

Mandiant has identified threat actors targeting virtualization platforms to achieve persistence and large-scale data destruction. Enforcing strict security controls through validated baselines is essential to disrupt these objectives. A defense-in-depth strategy that seeks to align virtual platform configurations with the official configuration and hardening guidance from the vendor can help provide a known security baseline for virtualization platforms.

A strong hardening posture begins with the hardware root of trust by activating UEFI Secure Boot and enabling the execInstalledOnly setting in conjunction with a supported TPM 2.0 chip. Secure Boot enforces kernel-level signature verification for all binaries while the TPM provides hardware-backed integrity checking of the boot process. This combination can prevent a threat actor from maintaining persistence through the addition of malicious installation bundles specific to the virtualization product or unauthorized drivers that would otherwise compromise the integrity of the hypervisor.

Going beyond host-based security controls, organizations should isolate the management plane to help prevent lateral movement from the general corporate environment. This can be achieved by restricting virtualization platform management traffic to a dedicated and firewalled network segment only accessible through hardened Privileged Access Workstations (PAW). To further minimize the attack surface, management shells and remote access protocols such as SSH should be disabled by default across all hosts. Direct remote access to the hypervisor should only be granted through a monitored "Break Glass" procedure for emergency troubleshooting. Infrastructure teams should partner with their security teams to ensure that high-priority detections exist for the specific break glass procedures they use.

Finally, protecting the most critical workloads requires the implementation of VM-level encryption for Tier-0 assets such as Domain Controllers. Encrypting these virtual machines can prevent threat actors from stealing virtual disks to steal sensitive data.

Backup Protection

Threat actors frequently prioritize the destruction of recovery infrastructure to eliminate an organization's ability to restore systems. To counter this, backup environments should be decoupled from the production virtualization environment and isolated from corporate identity stores. Integrating these platforms with the corporate Active Directory environment introduces a risk where an adversary can leverage a single compromised credential to delete both production and recovery data. Organizations should apply the same identity standards to backup infrastructure as they do to the virtualization platform, including phishing-resistant MFA and dedicated, non-AD integrated identity sources.

Resilience against data destruction requires immutable storage, such as Write-Once-Read-Many technology, to ensure backup repositories cannot be modified or encrypted by privileged accounts. Maintaining offsite and air-gapped copies of backup resources can provide a definitive barrier against network-based attacks. While this separation can introduce operational complexity, organizations should implement regular restoration testing from these isolated environments to verify that backups will function as intended during a high-pressure recovery event.

The Cascading Impact of Third-Party SaaS Compromises

The modern enterprise security landscape has shifted from the traditional network perimeter to a complex ecosystem of interconnected infrastructure and Software as a Service (SaaS) platforms. These integrated platforms play a vital role in identity management, workforce collaboration, and streamlining internal processes, among other applications. However, the interconnectedness of SaaS platforms with enterprise cloud infrastructure has introduced the risk of threat actors targeting identities used by SaaS integrations as a means to pivot into other parts of the environment. These incidents often occur through a combination of the exploitation of third-party integrations, over-permissive identities, and misconfigurations. Because of the deeply integrated nature of cloud and SaaS, a compromise on a single SaaS application can allow threat actors to move laterally into other areas of the environment with ease. In these scenarios, the failure of a singular trusted component triggers a chain reaction across the entire enterprise. As organizations accelerate towards cloud-first infrastructure, oversights such as unmonitored OAuth grants and broad API permissions increase the risk of a single compromised token causing a materially significant incident.

In 2025, a strategic shift emerged as threat actors bypassed traditional defenses like firewalls and multi-factor authentication (MFA) by abusing Non-Human Identities (NHIs) and stolen secrets such as OAuth and refresh tokens. By targeting vendor platforms that act as a centralized "source of truth" for identity integrations, adversaries can harvest legitimate, pre-authorized tokens to compromise downstream environments. This effectively transforms a single-vendor breach into a large-scale supply chain attack, where the stolen tokens act as reusable keys to the customers' sensitive data stores.

Lateral Movement Across Boundaries

One form of SaaS integration compromise that Mandiant encountered in 2025 involved threat actors compromising code repositories by gaining access to Personal Access Tokens. The threat actors then used automated scanning tools to search the code repositories for hardcoded integration tokens. This process often begins with a compromise on a SaaS developer's workstation or Source Code Management (SCM) account. Long-lived tokens found on the workstation or in the SCM account could allow the threat actor to pivot to the production cloud infrastructure. In some instances, access to cloud infrastructure could reveal OAuth refresh tokens for customer integrations, turning a single compromised repository into a potential supply chain event.

Observed Threat Actors and Trends in 2025

Throughout 2025, attackers exploited systemic weaknesses in how organizations manage SaaS interconnectivity and permissions for initial access. A primary driver of this was OAuth token compromise as some threat groups prioritized the theft of session cookies and refresh tokens over brute-forcing passwords. Because these tokens often remain valid post-logout, attackers can hijack sessions immediately without triggering MFA alerts. This risk is compounded by widespread SaaS misconfigurations, specifically those related to identities and secret management. Mandiant identified instances in which integrations that required only minimal access were provisioned API keys with broad privileges. For example, an integration intended only for log reading could be misconfigured with rights such as the ability to edit or delete storage buckets. Another common finding included the existence of API keys with wildcard permissions. Due to the prevalence and potential impact of these vulnerabilities, some threat actors have started to focus on stealing NHIs for third-party integrations from SaaS vendors. In 2025, Mandiant identified evidence of threat actors using NHIs, which were stolen during previous engagements.

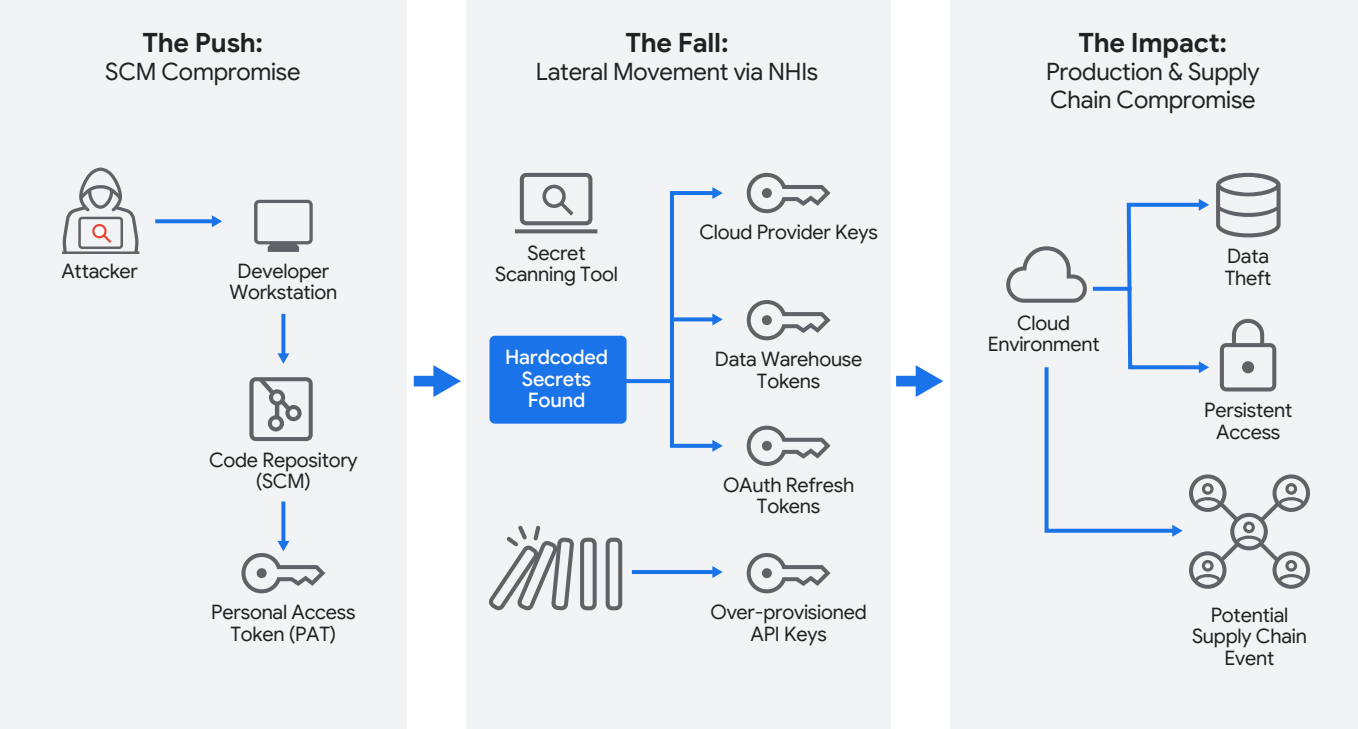

Figure 4: The SaaS Domino Effect Attack Chain

Mandiant has identified distinct threat clusters driving this shift, most commonly financially motivated actors. On the human-centric side of the spectrum, groups such as the financially motivated threat cluster UNC3944 prioritize aggressive social engineering over technical exploitation for initial access. These actors frequently target IT help desks, utilizing voice-based phishing (vishing) to impersonate employees and convince support staff to bypass MFA or enroll attacker-controlled devices. In some cases, following the compromise of a user account in a SaaS platform through similar human-centric techniques to those used by UNC3944, threat actors performed credential mining in the target environments for highly privileged keys that would allow them to compromise further downstream customer environments.

Groups like UNC6395 and UNC6564 demonstrated a focus on identity and SaaS integrations in 2025. UNC6395, a threat group Google Threat Intelligence Group (GTIG) began tracking in 2025, compromised a SaaS vendor and obtained a large number of SaaS integration tokens, which they used to perform mass data exports from downstream organizations with native command-line interfaces. They then ran automated scanners on the exported data. Their goal was to locate hardcoded keys, tokens, and Personal Access Tokens, which could be used for further downstream compromise. UNC6564, a financially motivated cluster, exemplified "Living off the Cloud" tactics, in which attackers rely on native cloud infrastructure and tools. This group operated from dedicated cloud infrastructure to blend in with legitimate traffic and abused previously compromised application tokens to gain access to target environments, then stole sensitive data.

Prevention

Preventing a third-party SaaS application compromise from becoming a materially significant event involves a defense-in-depth approach similar to a traditional enterprise environment, but with specific considerations due to the nature of SaaS applications. There are two threat profiles to keep in mind when designing prevention strategies for SaaS: third-party compromise, which can include supply chain attacks, and the compromise of internal human and non-human identities linked to SaaS applications. The methods for both remediating and proactively preventing SaaS compromise from either threat profile are similar. Security controls need to shift from static network blocking to dynamic identity

verification with identity being the common denominator among most SaaS attacks. A proactive prevention strategy incorporating these elements should be built on four pillars: knowing the environment, strategy and policy, identity and access management hardening, and continuous lifecycle management and detection.

Environment Inventory and Controls

Effective SaaS security management begins with comprehensive visibility and an understanding of the interconnected application landscape. A fundamental step is to maintain a complete inventory of all SaaS applications used across the enterprise. This inventory should go beyond just listing applications; it must include business owners, the types of data processed or stored, and all integration points. Organizations often overlook SaaS applications in traditional asset management, creating significant visibility gaps. Additionally, organizations must understand how data flows between integrated systems in order to identify high-risk pathways and data aggregation points, and to implement Data Loss Prevention (DLP) controls where sensitive data may be leaving the administrative boundaries of the environment.

Furthermore, proactive secret discovery and management are essential. This involves inventorying all secrets, including API keys, OAuth tokens, and service account credentials used for SaaS integrations. It is critical to identify where these secrets are stored, their dates of expiration, and who has access. Automated tools can scan code repositories, configuration files, and cloud environments for hardcoded or insecurely stored secrets. To help prevent unintentional inclusion of credentials in private or public code repositories, pre-commit hooks can be configured to identify strings that match common formats and halt commits to a repository. Enforcing pre-commit hooks on every repository, even those intended for personal use, helps reduce the risk of unintentional credential exposure.

To aid in managing this complex environment, organizations can deploy Cloud Access Security Brokers for visibility and control over SaaS application usage, and SaaS Security Posture Management tools to assess and remediate misconfigurations continuously across the SaaS estate. Given the diverse configuration options of each SaaS application, centralized management and automated posture checks are crucial. Finally, organizations should inventory and assess the logging and monitoring capabilities of each SaaS application, identify gaps in telemetry, and ensure critical logs are being ingested into a centralized security monitoring platform.

Strategy and Policy

A strong governance framework is essential to manage the risks associated with SaaS adoption and integration. This begins with a robust Third-Party Risk Management program to vet SaaS vendors before onboarding and on an ongoing basis. Key evaluation criteria include the vendor's security certifications, data handling practices, support for Single Sign-On, granular audit logging, and secure development practices. Organizations should mandate necessary security features, which are often not standard, as part of the procurement process.

Clear policies must also be established and enforced. Enforcing a policy requiring all SaaS applications to be integrated with the corporate Identity Provider (IdP) using secure federation protocols centralizes identity management and reduces reliance on application-specific credentials. Policies for the secure storage, rotation, and access control of all secrets should mandate the use of approved secrets management solutions and prohibit the use of hardcoded secrets. Organizations can create a clear hierarchy of acceptable use policies, security policies, standards, and procedures that explicitly address SaaS security, covering data handling, identity and access management, incident response, and vendor management. Lastly, it is critical to maintain a risk register to document and track known risks associated with SaaS applications, integrations, and vendors, reviewing and prioritizing risks for remediation regularly.

Strong policies governing end-user permissions and capabilities are also a critical component of SaaS security. Administrators can disable the ability for end users to consent to unverified third-party applications, ensuring that only vetted applications can obtain persistent tokens and effectively blocking the path for malicious OAuth applications to be granted those tokens.

Identity and Access Management Hardening

Given that identity is part of the new perimeter, hardening identity and access controls is paramount. Enforcing the principle of least privilege for all integrations can help preclude over-provisioned API keys and OAuth tokens. This includes avoiding wildcard permissions and granting only the specific scopes required for the integration to function, as well as assigning unique identities and credentials for each integration to limit the potential impact of a compromise of a single integration. Where possible, transitioning from static API keys and long-lived tokens to OpenID Connect (OIDC) federation for integrations between SaaS platforms and the IdP reduces the attack surface associated with stolen secrets. To minimize standing privileges for both human and non-human identities, organizations can implement Just-in-Time access mechanisms, granting elevated or sensitive permissions only for the duration needed to complete a specific task. Sender-constrained tokens can mitigate token theft and replay attacks through mechanisms like Demonstrated Proof-of-Possession or enforcing device compliance and context-aware access policies, ensuring tokens are only valid from trusted devices and locations.

Continuous Lifecycle Management and Detection

SaaS security is not a one-time setup; it requires continuous monitoring, maintenance, and adaptation. Organizations can implement automated processes for rotating all secrets, including API keys, OAuth refresh tokens, and service account credentials. Enforcing short lifespans for access tokens and browser sessions minimizes the window of opportunity for attackers leveraging stolen cookies or tokens; this is particularly critical as infostealer logs and active session cookies are traded on underground markets, where the value of a credential drops significantly the moment it expires or is rotated. Administrators can conduct periodic access reviews and recertification campaigns for both human and non-human identities to ensure permissions are still required and appropriate, automatically revoking unused permissions after a defined period. Similarly, implementing audit processes to detect and disable dormant user accounts as well as unused integrations helps reduce the attack surface available to threat actors. Organizations can also integrate targeted dark web monitoring to identify when specific corporate credentials or session tokens are being sold or shared as an early warning to trigger immediate rotation. By implementing multi-layered prevention strategies, organizations can significantly reduce the risk of a single point of failure within their SaaS ecosystem.

Outlook

Organizations are no longer defending a static fortress; instead, they are securing an interconnected web of trust relationships where the perimeter is defined not just by firewalls, but also by the entire identity and access management layer. To adapt to this shift, organizations must move beyond the assumption that any authenticated human user, service account, or third-party integration is inherently safe. Strategies for securing these assets must change from static prevention to continuous verification. This requires a cultural shift towards an architecture that forces IT teams to treat identity verification as a frontline defense. By recognizing that a risk accepted in a third-party application deployed into an interconnected environment is also a risk in the production core, organizations reframe their view of identity and access controls. Enforcing policies that align to the principle of least privilege, and adopting a "Never Trust, Always Verify" mindset allows organizations to get ahead of threat actors that target third-party SaaS integrations and turn potential systemic failure into a contained, manageable event.

Systematic Exploitation of Edge and Core Network Devices

Over the past several years, Mandiant has seen an increase of malicious actors targeting edge and core network devices. Attackers have taken advantage of the fact that many of these devices are not able to run enterprise security tooling that provides increased visibility into threat actor actions and can interdict and prevent malicious activity.

At the same time, Google Threat Intelligence Group (GTIG) has observed that the mean time to exploit (TTE) for vulnerabilities has continually decreased from 63 days in 2018 to -1 day in 2024 and further downward to an estimated -7 days in 2025. A negative number indicates that exploitation of a vulnerability, on average, occurred before a patch was released. This trend is compounded by threat actors' evolving interest in security appliances and networking infrastructure. Because these assets provide critical inspection of data ingress and egress points, they grant significant visibility and permissions to those who gain access to them. Additionally, they often show signs of extended uptime periods, delays in patching, and can often be excluded from vulnerability management altogether. This has allowed threat actors the opportunity to exploit n-day vulnerabilities in these devices repeatedly without needing to identify novel vulnerabilities.

Attackers have long used edge devices as a launch point into targeted environments. Historically, systems with traditional operating systems such as Windows, Linux, and macOS have been the primary operating space attackers have used to complete their mission objectives. However, in recent years, Mandiant has observed attackers shift tactics, increasingly performing more phases of the attack lifecycle from edge and core network devices. Many of these devices offer built-in administrative functionality. Threat actors have leveraged this capability in a number of ways, including to facilitate long term access, perform reconnaissance, move laterally through an environment, escalate privileges, and perform data collection. Mandiant has identified threat actors, in some cases, pivot to collecting the data that traverses management planes or resides on core network devices themselves, as opposed to obtaining it from traditional sources such as workstations, servers or databases. This method of data collection has allowed these actors to further avoid security monitoring and remain undetected for long periods of time.

While the exploitation of edge devices and core network devices is not a new phenomenon, the abuse of these devices has steadily increased as the TTE continues to decrease. Investigating attacker activity performed on these devices presents a significant challenge for security teams. A lack of asset inventories and sub-standard telemetry originating from the devices themselves results in added difficulty when security teams attempt to assess evidence of compromise. Edge and core network devices are designed to have minimal operating systems to maximize computing power. As a result, traditional file system forensics, log analysis, and memory forensics can be limited, as onboard storage for these devices is typically limited to storing configuration files and required binaries. All of these limitations compound so that when security teams do identify suspicious activity on these assets, there are often limited forensic artifacts to confirm an attacker's presence or that a vulnerability was successfully exploited.

In 2025, Mandiant investigated a diverse group of attackers targeting edge devices. The espionage actor, UNC5807, used their access to core network devices to collect data for intelligence purposes. Another intelligence-focused actor, UNC5221, repeatedly targeted edge devices, demonstrating an extensive knowledge of security appliances and using both zero-day and n-day flaws in order to maintain long-term stealthy access to organizations.

Real-World Examples

UNC5807

The feature set of many core and edge network appliances has grown beyond relatively simple packet routing and switching devices. As the tools have grown in complexity, that same complexity has often come at the cost of security instrumentation, high friction, and costly logging infrastructure. Threat actors have started to target this gap between feature sets and a lagging security posture by targeting these platforms as a primary aspect of their campaigns within an environment.

UNC5807, a sophisticated PRC-nexus espionage operator, has been observed targeting the telecommunications sector. UNC5807's activity is generally consistent with the threat actor discussed in public reporting as "Salt Typhoon". In multiple instances, UNC5807's objectives appeared to align with the collection of highly sensitive data for intelligence-gathering purposes. Mandiant identified evidence that the threat group targeted internet-accessible network infrastructure within a variety of organizations and relied primarily on the feature sets inherent to those platforms to scope the environment before moving deeper into the network. As they moved through the environment, UNC5807 followed a demonstrable playbook in which high-impact actions were limited to areas of the environment least likely to raise alarm.

UNC5807 relied heavily on application subshells built into the platforms they compromised. These subshells exposed features that are commonly restricted for normal users but which enhanced the visibility and reach UNC5807 had into the targeted environment. In order to maintain long-term multi-layered access to compromised assets, the threat group created local admin accounts on the appliances and enabled services such as SSH on non-standard ports.

The relatively long uptime common to core and edge network appliances provided long-term stability for UNC5807's persistence, as well as for data collection activities. UNC5807 used the packet-capturing functionality built into the platforms to collect copies of live traffic that transited the devices. By analyzing the captured data, UNC5807 identified passwords from cleartext network protocols and, using the various functions built into the subshell, leveraged those credentials to move through the network. Once the threat actors identified a system through which sensitive data was transmitted, they used device-native utilities to capture live traffic. This data was staged on systems which supported common file transfer protocols and was subsequently moved to devices in the edge network where traffic and security monitoring were less comprehensive. Once staged in the edge network, UNC5807 uploaded the data to threat actor-controlled infrastructure.

UNC5807's successes in compromising and using edge and core network appliances highlight the need for a comprehensive threat model that can be applied to common user devices and relatively obscure corporate infrastructure devices alike. The ability to launch tooling that provides admin-level capabilities on systems which are rarely patched and power cycled—much less actively monitored for access—represents an enticing feature set for threat actors.

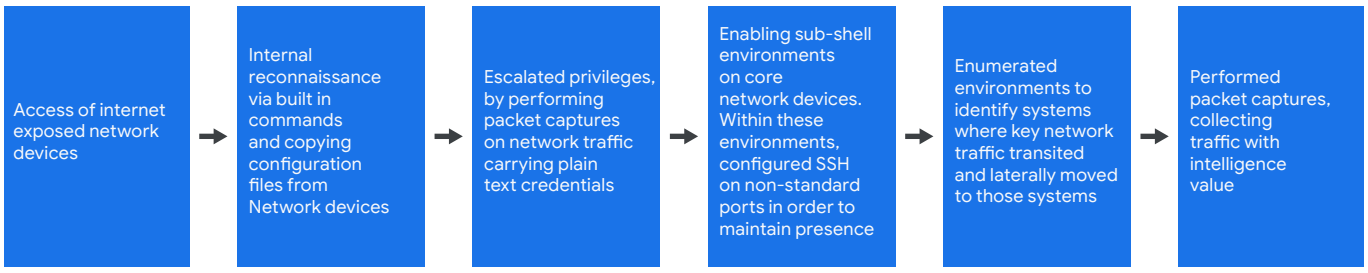

Figure 5: UNC5807 Attack Flow

UNC5221

UNC5221 is an advanced PRC-nexus espionage actor and has been observed targeting a range of different industries. While UNC5221 has been used synonymously with the threat actor publicly reported as "Silk Typhoon," GTIG does not currently consider the two clusters to be the same. UNC5221's operations are characterized by the use of custom malware toolsets designed for stealth and persistence. In 2025, UNC5221 demonstrated a relentless focus on edge devices by exploiting two unique zero-day vulnerabilities in the same VPN product during a three-month period. This campaign highlights a strategic drive to maintain access to these critical gateways despite vendor remediation efforts. Upon gaining access to a targeted environment, UNC5221 prioritized stealth and evasion, deploying specialized malware such as TRAILBLAZE and BRUSHFIRE directly to the appliance. TRAILBLAZE is an in-memory only dropper that injects a hook into the web process. BRUSHFIRE is a passive backdoor that acts as an SSL_read hook. It first executes the original SSL_read function and checks to see if the returned data begins with a specific string. If the data begins with the defined marker string, it will XOR decrypt then execute shellcode contained in the data. Consistent with their previous campaigns targeting VPN edge devices, UNC5221 relied on non-persistent tactics and aggressive on-disk cleanup to remove evidence of exploitation and minimize the risk of discovery.

Using the pivot point provided by the compromised VPN edge devices, UNC5221 utilized legitimate credentials harvested from the appliance's own configuration files to authenticate to the VPN service and move laterally through the environment. The group followed a methodical approach to privilege escalation, including targeting backup and infrastructure management systems, as well as analyzing administrative history log files and internal documentation repositories for plaintext credentials. In one case, UNC5221 accessed a backup infrastructure device where they exported a privileged authentication management system backup and uploaded it to a cloud-based storage service. The threat actors were able to analyze the backup and gain access to privileged credentials outside of the targeted environment.

To solidify long-term access, UNC5221 moved laterally into the virtualization and management layers of the network where they compromised virtualization infrastructure and support management appliances. These systems provided a privileged vantage point for internal reconnaissance, allowing the actor to map virtual machines and scan the network topology. To maintain command and control (C2) without raising alarms, the group deployed lightweight backdoors which tunneled communications over HTTPS through legitimate third-party services. By using trusted external platforms for C2 traffic, UNC5221 was able to bypass standard network security filtering while maintaining persistent control over the compromised infrastructure devices.

UNC5221 has demonstrated a pattern of maintaining long-term stealthy access to environments by targeting the visibility gap. By focusing their efforts on devices that do not support traditional endpoint detection and response (EDR) tools, as well as avoiding actions that are likely to raise alarms, UNC5221 dramatically reduced the opportunities in which defenders could identify their activity. UNC5221's success in maintaining long-term access using these methods emphasizes the need for organizations to improve visibility and logging for non-traditional assets, as these devices are increasingly becoming attractive targets for threat actors.

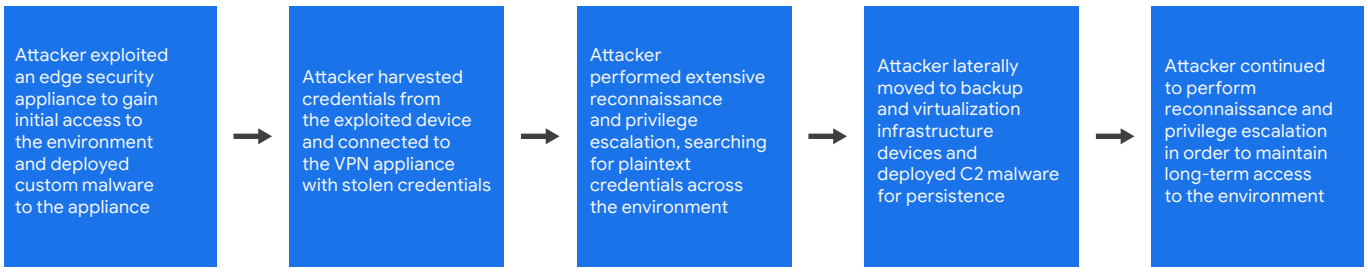

Figure 6: UNC5221 Attack Flow

Hardening, Visibility

The challenge of detecting threats originating from edge and core network devices primarily stems from three factors: limited visibility, inadequate documentation, and strict business continuity requirements. The general lack of visibility into actions performed on or from networking devices, beyond standard traffic logging, represents a significant friction point for successful detection and response efforts. Common enterprise logging configurations provide a foundational security and business continuity layer. However, by focusing on user activity on endpoints as the primary sources of data to forward to a centralized location, security teams frequently exclude critical log sources from core and edge network devices. The large volume of this data relative to its perceived security value also makes it difficult for analysts to sift through and identify threat actor activity. Typically, log data for network devices is focused around the Network, Transport and Session layers. Application logging, such as accounting logs, are usually stored locally, if at all. However, these logs are among the strongest data sets available for a security team to identify anomalies during an investigation into possible network device abuse.

A successful logging strategy should balance resource efficiency with the preservation of critical evidence. This relies on two main pillars: aligning logging objectives with the organization's specific threat model, and integrating logging as a core requirement within the system development lifecycle for all new and acquired systems. Logging requirements should be reviewed regularly to ensure alignment with organizational goals; retention times should be adjusted against the Mean Time to Detect to enable shorter and more cost-effective retention periods. In general, Mandiant recommends minimum retention periods of 1 year for administrative activity logs, 90 days for security-related events, and 30 days for informational logs used to contextualize other events.

The specialized nature of edge and core network device configurations often results in a knowledge gap for security personnel, who typically focus on endpoint analysis, packet captures, and basic TCP/UDP protocols. Security teams should maintain detailed playbooks outlining the network architecture, including diagrams, device administrators, IP addresses, and hostnames. These playbooks are vital for technical responders to follow critical steps and for executives to make sound decisions during a security event.

Organizations should also commit to regularly testing and updating these response playbooks through tabletop exercises. Additionally, it is prudent to understand, before an incident, precisely how and what forensic collections could be taken from network devices. Given that network devices are rarely altered or updated due to business criticality, the collection process requires robust planning to ensure business continuity. This pre-incident planning should define whether external support, such as a vendor, is required, and what the potential limitations of the forensic collections supported mean for analysis. During collection, preserving the integrity of evidence is paramount; forensic artifacts may be lost if a system is powered off. Organizations should establish documented procedures for creating, storing, and transferring forensic images, including retaining the cryptographic hash for validation and meticulously recording all actions taken to maintain chain of custody.

Due to common delays in patching and power cycling of edge and core network devices, the processes often require more robust planning to ensure the network can either continue to function or prepare for downtime. Additionally, vulnerability management is essential not only for identifying immediate risks but also for serving as a secondary check for asset management by discovering uncatalogued systems. Program success is dependent on timely remediation. Clearly defined oversight mechanisms and ownership for mitigation ensure accountability, as unaddressed vulnerabilities remain a primary target for unauthorized access. To enhance vulnerability management processes, organizations should leverage their existing vulnerability management solution to conduct comprehensive network discovery scans that include fingerprinting publicly available devices and IP address space from unauthenticated sources. While exploiting zero-day and n-day vulnerabilities are the most common initial access vectors, these devices are also susceptible to abuse of legitimate protocols such as publicly available FTP and SSH services, using stolen credentials.

Organizations should establish a formalized, staggered assessment schedule to ensure total visibility without overtaxing network resources. Scans should be deployed in phases until enterprise-wide coverage is achieved. The objective is to transition into a consistent cadence where every endpoint is assessed for vulnerabilities at least twice monthly. Vulnerability data should be analyzed through a business context lens, prioritizing remediation based on the asset's value and its impact on business operations. To drive this process, the

vulnerability management team should distribute targeted reports to stakeholders that include a clear description of the identified risk, a list of affected endpoints, the severity level specific to each asset, identification of deviations from internal security policies, and specific recommendations for remediation.

Furthermore, security teams should be included in general architecture discussions to develop relationships between teams, ensure proper telemetry is collected, and to establish a clear understanding of where devices fit into the network. This can help ensure that these critical devices are not identified as a problem only after a threat has moved deeper into the environment.

Outlook

Security monitoring and controls on traditional assets continue to gain greater adoption and higher effectiveness at identifying and stopping threats. At the same time, edge devices have seen an increase in the number of vulnerabilities identified. While edge and core network devices are not novel targets for sophisticated actors, Mandiant expects these devices to become increasingly high priority targets for actors of all levels of sophistication. The challenges in security instrumentation combined with the difficulty in remediating compromised edge and core network devices present a significant headwind for network defenders. Threat actors will likely shift their tactics toward performing more of the attack lifecycle on this class of device and away from traditional assets. Furthermore, Mandiant expects exploitation of network devices will continue to be an enticing initial access vector for threat actors, as they look for remaining safe havens in targeted environments.

Given the exposure, lack of visibility and critical location of these devices in most networks, organizations must take steps to identify the risks that unauthorized access to these systems pose. Organizations should work to identify these assets in their networks and implement robust patch management, logging, and access controls to limit the risk associated with compromise of these devices.

Conclusion

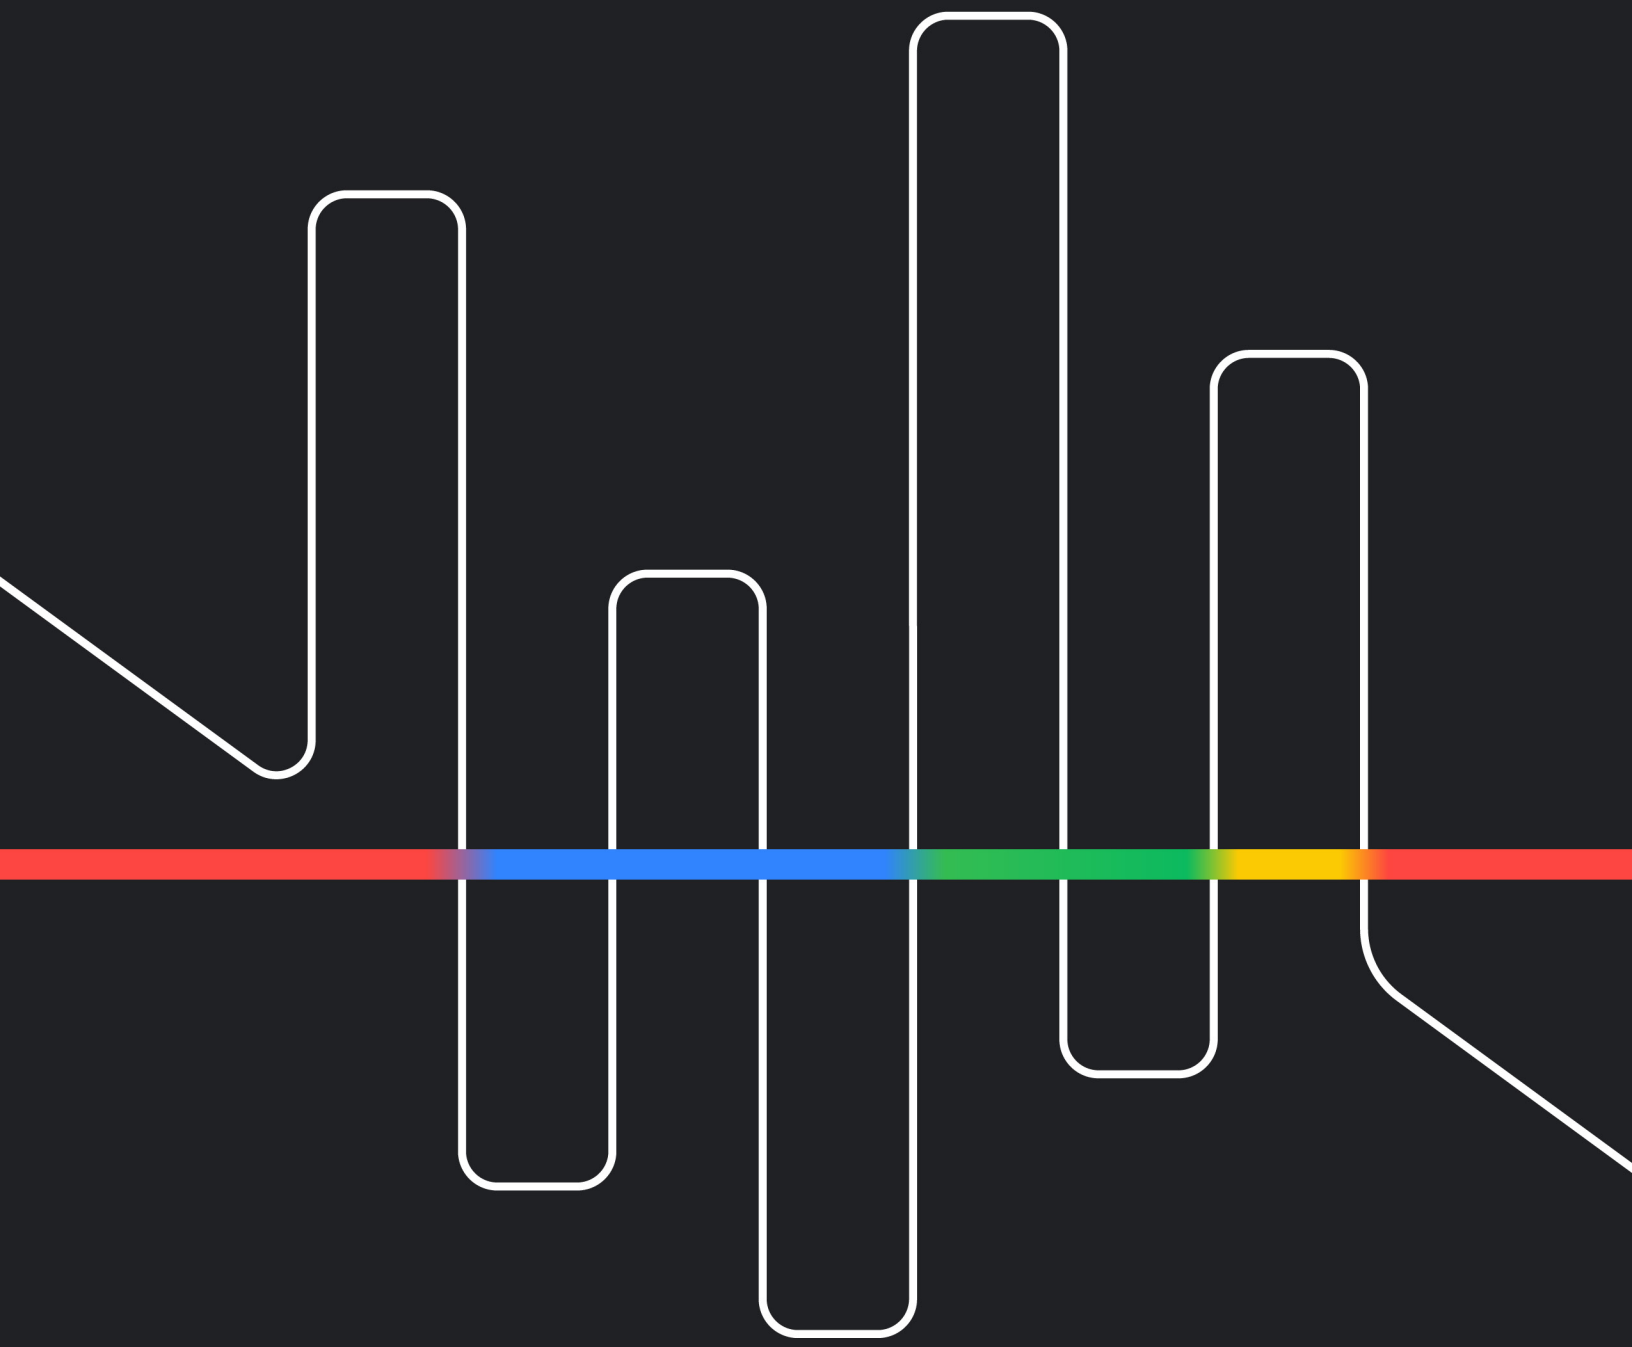

Our threat intelligence analysts have observed adversaries leveraging AI to accelerate attacks, shifting from mass, static email campaigns toward hyper-personalized, voice-based social engineering, and deploying malware capable of querying large language models mid-execution. Yet, despite these rapid technological advancements, the incidents investigated by Mandiant for M-Trends 2026 primarily stemmed from fundamental human and systemic failures.

We saw a stark divergence in adversary pacing. While cybercriminal groups optimized for immediate impact and deliberate recovery denial, systematically targeting backup infrastructure, identity services, and virtualization management planes, cyber espionage groups and DPRK IT workers optimized for extreme persistence. By operating from unmonitored edge devices and utilizing native network functionalities to evade detection, these stealthy actors drove the global median dwell time up to 14 days.

Defending against these tactics, techniques and procedures requires organizations to move at the speed of the adversary. A proactive, resilient defense must go beyond static tools; it necessitates continuous identity verification, rigorous protection of critical control planes, and expansive visibility across the entire ecosystem, including the virtualization layer and EDR-less network appliances. Security teams must be rigorously tested through realistic red team engagements that incorporate modern AI-driven tactics. Furthermore, organizations should conduct regular tabletop exercises to update incident response playbooks for modern extortion pipelines, and maintain a cyber incident response retainer to ensure immediate access to expert help before a minor alert becomes a catastrophic compromise.

Exploits (32%) remained the most common initial infection vector for the sixth consecutive year, followed closely by a significant surge in highly interactive voice phishing (11%). Foundational security practices must evolve to meet these specific realities. Organizations must aggressively manage internet-facing attack surfaces, strictly isolate external-facing web application servers, and pivot security awareness training beyond the inbox to address live, interactive social engineering.

Building a comprehensive security program now means securing developer and AI toolchains, treating identity as the ultimate perimeter, and ensuring recovery capabilities are fundamentally segmented from the production environment.

The Mandiant mission is to help keep every organization secure from cyber threats and confident in their readiness. Our annual M-Trends report, powered by the collective intelligence of frontline incident responders and the Google Threat Intelligence Group, is a core component of advancing that mission. We will continue to share our frontline knowledge to help defenders close critical visibility gaps, outmaneuver modern adversaries, and build true operational resilience.

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 front-line 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

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

Techniques Related to Mandiant Targeted Attack Lifecycle, 2025

☁ Indicates techniques in the Cloud matrix, introduced in ATT&CK v18.

Initial Reconnaissance

Reconnaissance

T1593: Search Open Websites/Domains	16.3%		
T1590: Gather Victim Network Information	4.4%	T1590.002: DNS	4.4%
T1592: Gather Victim Host Information	2.9%		
T1598: Phishing for Information	2.4%		
T1595: Active Scanning	1.0%	T1595.002: Vulnerability Scanning	0.5%
		T1595.001: Scanning IP Blocks	0.2%
		T1595.003: Wordlist Scanning	0.2%
T1589: Gather Victim Identity Information	0.2%	T1589.001: Credentials	0.2%

Resource Development

T1608: Stage Capabilities	11.7%	T1608.001: Upload Malware	3.9%
		T1608.003: Install Digital Certificate	2.9%
		T1608.005: Link Target	2.2%
		T1608.002: Upload Tool	1.9%
		T1608.006: SEO Poisoning	0.5%
T1588: Obtain Capabilities	8.5%	T1588.003: Code Signing Certificates	8.3%
		T1588.007: Artificial Intelligence	0.2%
T1583: Acquire Infrastructure	6.8%	T1583.003: Virtual Private Server	6.8%
T1584: Compromise Infrastructure	3.2%		
T1587: Develop Capabilities	0.5%	T1587.002: Code Signing Certificates	0.2%
		T1587.003: Digital Certificates	0.2%

Initial Compromise

Initial Access

T1190: Exploit Public-Facing Application ☁	27.7%		
T1133: External Remote Services	23.8%		
T1078: Valid Accounts ☁	17%	T1078.004: Cloud Accounts ☁	9.5%
T1566: Phishing ☁	14.1%	T1566.004: Spearphishing Voice ☁	6.1%
		T1566.002: Spearphishing Link ☁	3.2%
		T1566.003: Spearphishing via Service	2.7%
		T1566.001: Spearphishing Attachment	1.0%
T1189: Drive-by Compromise ☁	6.3%		
T1195: Supply Chain Compromise ☁	1.7%	T1195.002: Compromise Software Supply Chain	1.5%
		T1195.003: Compromise Hardware Supply Chain	0.2%
T1200: Hardware Additions	1.0%		
T1091: Replication Through Removable Media	0.5%		
T1199: Trusted Relationship ☁	0.5%		
T1659: Content Injection	0.2%		

Initial Compromise

Credential Access

T1003: OS Credential Dumping	13.6%	T1003.003: NTDS	4.4%
		T1003.008: /etc/passwd and /etc/shadow	3.2%
		T1003.001: LSASS Memory	3.2%
		T1003.002: Security Account Manager	2.4%
		T1003.006: DCSync	1.7%
		T1003.004: LSA Secrets	0.2%
T1552: Unsecured Credentials ☁	10.4%	T1552.002: Credentials in Registry	2.7%
		T1552.001: Credentials In Files ☁	1.7%
		T1552.003: Bash History	1.5%
		T1552.007: Container API	1.2%
		T1552.004: Private Keys	1.0%
		T1552.006: Group Policy Preferences	0.2%
		T1552.005: Cloud Instance Metadata API ☁	0.2%
T1555: Credentials from Password Stores ☁	6.1%	T1555.006: Cloud Secrets Management Stores ☁	2.4%
		T1555.003: Credentials from Web Browsers	2.2%
		T1555.005: Password Managers	1.0%
		T1555.004: Windows Credential Manager	0.2%
T1111: Multi-Factor Authentication Interception	4.9%		
T1110: Brute Force ☁	4.1%	T1110.001: Password Guessing ☁	1.5%
		T1110.004: Credential Stuffing ☁	1.0%
		T1110.003: Password Spraying ☁	0.7%
T1556: Modify Authentication Process ☁	2.7%	T1556.006: Multi-Factor Authentication ☁	1.5%
		T1556.003: Pluggable Authentication Modules	0.2%
		T1556.002: Password Filter DLL	0.2%
T1558: Steal or Forge Kerberos Tickets	2.7%	T1558.003: Kerberoasting	1.9%
		T1558.001: Golden Ticket	0.2%
		T1558.004: AS-REP Roasting	0.2%
		T1558.002: Silver Ticket	0.2%
T1040: Network Sniffing ☁	1.0%		
T1187: Forced Authentication	0.7%		
T1528: Steal Application Access Token ☁	0.2%		
T1539: Steal Web Session Cookie ☁	0.2%		
T1649: Steal or Forge Authentication Certificates ☁	0.2%		

Establish Foothold

Execution

T1059: Command and Scripting Interpreter ☁	45.9%	T1059.003: Windows Command Shell	26.2%
		T1059.001: PowerShell	24%
		T1059.007: JavaScript	5.6%
		T1059.004: Unix Shell	5.3%
		T1059.006: Python	4.6%
		T1059.005: Visual Basic	3.2%
		T1059.012: Hypervisor CLI	2.2%
		T1059.010: AutoHotKey & AutoIT	0.5%
		T1059.009: Cloud API ☁	0.5%
T1204: User Execution ☁	19.9%	T1204.002: Malicious File	16.5%
		T1204.001: Malicious Link	2.2%
		T1204.004: Malicious Copy and Paste	1.2%
T1569: System Services	15.0%	T1569.002: Service Execution	14.3%
		T1569.003: Systemctl	0.5%
T1053: Scheduled Task/Job	11.2%	T1053.005: Scheduled Task	9.0%
		T1053.003: Cron	1.7%
T1047: Windows Management Instrumentation	7.0%		
T1129: Shared Modules	2.7%		
T1559: Inter-Process Communication	1.9%		
T1651: Cloud Administration Command ☁	0.5%		
T1203: Exploitation for Client Execution	0.2%		
T1609: Container Administration Command	0.2%		
T1610: Deploy Container	0.2%		

Establish Foothold

Persistence

T1133: External Remote Services	23.8%		
T1098: Account Manipulation ☁	18.7%	T1098.007: Additional Local or Domain Groups	7.8%
		T1098.005: Device Registration ☁	3.6%
		T1098.001: Additional Cloud Credentials ☁	1.0%
		T1098.004: SSH Authorized Keys ☁	0.5%
		T1098.003: Additional Cloud Roles ☁	0.2%
T1078: Valid Accounts ☁	17.0%	T1078.004: Cloud Accounts ☁	9.5%
T1543: Create or Modify System Process	16.7%	T1543.003: Windows Service	8.3%
		T1543.002: Systemd Service	1.9%
		T1543.004: Launch Daemon	0.2%
T1505: Server Software Component	14.8%	T1505.003: Web Shell	14.8%
T1053: Scheduled Task/Job	11.2%	T1053.005: Scheduled Task	9.0%
		T1053.003: Cron	1.7%
T1136: Create Account ☁	10.2%	T1136.001: Local Account	3.2%
		T1136.003: Cloud Account ☁	1.5%
		T1136.002: Domain Account	1.0%
T1547: Boot or Logon Autostart Execution	8.3%	T1547.001: Registry Run Keys / Startup Folder	7.5%
		T1547.009: Shortcut Modification	1.0%
		T1547.013: XDG Autostart Entries	0.5%
		T1547.005: Security Support Provider	0.2%
T1574: Hijack Execution Flow	7.3%	T1574.011: Services Registry Permissions Weakness	6.8%
		T1574.001: DLL	0.5%
		T1574.008: Path Interception by Search Order Hijacking	0.2%
T1546: Event Triggered Execution ☁	6.1%	T1546.003: Windows Management Instrumentation Event Subscription	4.1%
		T1546.015: Component Object Model Hijacking	0.7%
		T1546.007: Netsh Helper DLL	0.5%
		T1546.010: AppInit DLLs	0.2%
		T1546.004: Unix Shell Configuration Modification	0.2%
		T1546.001: Change Default File Association	0.2%
T1671: Cloud Application Integration ☁	4.6%		
T1556: Modify Authentication Process	1.5%	T1556.006: Multi-Factor Authentication ☁	1.5%
T1037: Boot or Logon Initialization Scripts	0.5%	T1037.001: Logon Script (Windows)	0.2%
		T1037.004: RC Scripts	0.2%
T1176: Software Extensions	0.5%	T1176.001: Browser Extensions	0.2%
T1137: Office Application Startup ☁	0.2%	T1137.005: Outlook Rules ☁	0.2%
T1554: Compromise Host Software Binary	0.2%		

Escalate Privileges

Credential Access

T1003: OS Credential Dumping	13.6%	T1003.003: NTDS	4.4%
		T1003.008: /etc/passwd and /etc/shadow	3.2%
		T1003.001: LSASS Memory	3.2%
		T1003.002: Security Account Manager	2.4%
		T1003.006: DCSync	1.7%
		T1003.004: LSA Secrets	0.2%
T1552: Unsecured Credentials ☁	10.4%	T1552.002: Credentials in Registry	2.7%
		T1552.001: Credentials In Files ☁	1.7%
		T1552.003: Bash History	1.5%
		T1552.007: Container API	1.2%
		T1552.004: Private Keys	1.0%
		T1552.006: Group Policy Preferences	0.2%
		T1552.005: Cloud Instance Metadata API ☁	0.2%
T1555: Credentials from Password Stores ☁	6.1%	T1555.006: Cloud Secrets Management Stores ☁	2.4%
		T1555.003: Credentials from Web Browsers	2.2%
		T1555.005: Password Managers	1.0%
T1111: Multi-Factor Authentication Interception	4.9%		
T1110: Brute Force ☁	4.1%	T1110.001: Password Guessing ☁	1.5%
		T1110.004: Credential Stuffing ☁	1.0%
		T1110.003: Password Spraying ☁	0.7%
T1556: Modify Authentication Process ☁	2.7%	T1556.002: Password Filter DLL	2.0%
		T1556.003: Pluggable Authentication Modules	2.0%
		T1556.006: Multi-Factor Authentication ☁	1.5%
T1558: Steal or Forge Kerberos Tickets	2.7%	T1558.003: Kerberoasting	1.9%
		T1558.001: Golden Ticket	0.2%
		T1558.004: AS-REP Roasting	0.2%
		T1558.002: Silver Ticket	0.2%
T1040: Network Sniffing ☁	1.0%		
T1187: Forced Authentication	0.7%		
T1528: Steal Application Access Token ☁	0.2%		
T1539: Steal Web Session Cookie ☁	0.2%		
T1649: Steal or Forge Authentication Certificates ☁	0.2%		

Escalate Privileges

Privilege Escalation

T1078: Valid Accounts ☁	17%	T1078.004: Cloud Accounts ☁	9.5%
T1543: Create or Modify System Process	16.7%	T1543.003: Windows Service	8.3%
		T1543.002: Systemd Service	1.9%
		T1543.004: Launch Daemon	0.2%
T1053: Scheduled Task/Job	11.2%	T1053.005: Scheduled Task	9.0%
		T1053.003: Cron	1.7%
T1547: Boot or Logon Autostart Execution	8.3%	T1547.001: Registry Run Keys / Startup Folder	7.5%
		T1547.009: Shortcut Modification	1.0%
		T1547.013: XDG Autostart Entries	0.5%
		T1547.005: Security Support Provider	0.2%
T1098: Account Manipulation	7.8%	T1098.007: Additional Local or Domain Groups	7.8%
T1574: Hijack Execution Flow	7.3%	T1574.011: Services Registry Permissions Weakness	6.8
		T1574.001: DLL	0.5%
		T1574.008: Path Interception by Search Order Hijacking	0.2%
T1546: Event Triggered Execution ☁	6.1%	T1546.003: Windows Management Instrumentation Event Subscription	4.1%
		T1546.015: Component Object Model Hijacking	0.7%
		T1546.007: Netsh Helper DLL	0.5%
		T1546.010: AppInit DLLs	0.2%
		T1546.004: Unix Shell Configuration Modification	0.2%
		T1546.001: Change Default File Association	0.2%
T1055: Process Injection	1.0%		
T1068: Exploitation for Privilege Escalation	1.0%		
T1548: Abuse Elevation Control Mechanism ☁	1.0%	T1548.002: Bypass User Account Control	0.5%
		T1548.003: Sudo and Sudo Caching	0.2%
		T1548.001: Setuid and Setgid	0.2%
T1134: Access Token Manipulation	0.7%	T1134.001: Token Impersonation/Theft	0.7%
T1484: Domain or Tenant Policy Modification ☁	0.7%	T1484.001: Group Policy Modification	0.7%
T1037: Boot or Logon Initialization Scripts	0.5%	T1037.001: Logon Script (Windows)	0.2%
		T1037.004: RC Scripts	0.2%

Internal Reconnaissance

Discovery

T1083: File and Directory Discovery	33.5%		
T1033: System Owner/User Discovery	24.0%		
T1087: Account Discovery ☁	20.9%	T1087.002: Domain Account	11.9%
		T1087.001: Local Account	10.7%
		T1087.004: Cloud Account ☁	1.9%
		T1087.003: Email Account ☁	0.2%
T1016: System Network Configuration Discovery	19.4%	T1016.001: Internet Connection Discovery	14.1%
		T1016.002: Wi-Fi Discovery	0.7%
T1082: System Information Discovery ☁	18.2%		
T1518: Software Discovery ☁	17.0%	T1518.001: Security Software Discovery	1.2%
T1057: Process Discovery	14.6%		
T1069: Permission Groups Discovery ☁	14.1%	T1069.002: Domain Groups	8.3%
		T1069.001: Local Groups	3.4%
		T1069.003: Cloud Groups ☁	1.7%
T1049: System Network Connections Discovery	10.0%		
T1482: Domain Trust Discovery	8.5%		
T1012: Query Registry	5.1%		
T1018: Remote System Discovery	4.9%		
T1007: System Service Discovery	3.9%		
T1580: Cloud Infrastructure Discovery ☁	3.4%		
T1619: Cloud Storage Object Discovery ☁	3.4%		
T1046: Network Service Discovery ☁	2.9%		
T1201: Password Policy Discovery ☁	1.7%		
T1614: System Location Discovery ☁	1.5%	T1614.001: System Language Discovery	1.5%
T1124: System Time Discovery	1.2%		
T1040: Network Sniffing ☁	1.0%		
T1654: Log Enumeration ☁	1.0%		
T1673: Virtual Machine Discovery	1.0%		
T1680: Local Storage Discovery	0.7%		
T1120: Peripheral Device Discovery	0.5%		
T1538: Cloud Service Dashboard	0.5%		
T1652: Device Driver Discovery ☁	0.5%		
T1217: Browser Information Discovery	0.2%		
T1526: Cloud Service Discovery ☁	0.2%		

Lateral Movement

Execution

T1059: Command and Scripting Interpreter ☁	45.9%	T1059.003: Windows Command Shell	26.2%
		T1059.001: PowerShell	24.0%
		T1059.007: JavaScript	5.6%
		T1059.004: Unix Shell	5.3%
		T1059.006: Python	4.6%
		T1059.005: Visual Basic	3.2%
		T1059.012: Hypervisor CLI	2.2%
		T1059.010: AutoHotKey & AutoIT	0.5%
		T1059.009: Cloud API ☁	0.5%
T1204: User Execution ☁	19.9%	T1204.002: Malicious File	16.5%
		T1204.001: Malicious Link	2.2%
		T1204.004: Malicious Copy and Paste	1.2%
T1569: System Services	15%	T1569.002: Service Execution	14.3%
		T1569.003: Systemctl	0.5%
T1053: Scheduled Task/Job	11.2%	T1053.005: Scheduled Task	9.0%
		T1053.003: Cron	1.7%
T1047: Windows Management Instrumentation	7.0%		
T1129: Shared Modules	2.7%		
T1559: Inter-Process Communication	1.9%		
T1651: Cloud Administration Command ☁	0.5%		
T1203: Exploitation for Client Execution	0.2%		
T1609: Container Administration Command	0.2%		
T1610: Deploy Container	0.2%		

Lateral Movement

Lateral Movement

T1021: Remote Services ☁	30.6%	T1021.001: Remote Desktop Protocol	22.1%
		T1021.002: SMB/Windows Admin Shares	19.4%
		T1021.004: SSH	12.9%
		T1021.006: Windows Remote Management	2.4%
		T1021.008: Direct Cloud VM Connections	0.5%
		T1021.003: Distributed Component Object Model	0.5%
		T1021.005: VNC	0.2%
T1550: Use Alternate Authentication Material ☁	6.3%	T1550.001: Application Access Token ☁	5.3%
		T1550.002: Pass the Hash	1.0%
T1570: Lateral Tool Transfer	1.0%		
T1091: Replication Through Removable Media	0.5%		
T1080: Taint Shared Content ☁	0.2%		

Maintain Presence

Defense Evasion

T1027: Obfuscated Files or Information	27.2%	T1027.015: Compression	5.6%
		T1027.002: Software Packing	4.4%
		T1027.010: Command Obfuscation	4.4%
		T1027.009: Embedded Payloads	2.2%
		T1027.004: Compile After Delivery	1.5%
		T1027.013: Encrypted/Encoded File	1.2%
		T1027.001: Binary Padding	0.2%
T1070: Indicator Removal ☁	24.8%	T1070.004: File Deletion	15.5%
		T1070.009: Clear Persistence	5.1%
		T1070.001: Clear Windows Event Logs	4.6%
		T1070.006: Timestomp	3.2%
		T1070.007: Clear Network Connection History and Configurations	1.9%
		T1070.002: Clear Linux or Mac System Logs	0.5%
		T1070.003: Clear Command History	0.5%
		T1070.010: Relocate Malware	0.5%
		T1070.008: Clear Mailbox Data ☁	0.2%
T1078: Valid Accounts ☁	17.0%	T1078.004: Cloud Accounts ☁	9.5%
T1562: Impair Defenses ☁	16.5%	T1562.001: Disable or Modify Tools ☁	10.2%
		T1562.004: Disable or Modify System Firewall	9.0%
		T1562.002: Disable Windows Event Logging	4.4%
		T1562.003: Impair Command History Logging	0.7%
		T1562.012: Disable or Modify Linux Audit System	0.2%
T1140: Deobfuscate/Decode Files or Information	15.3%		
T1564: Hide Artifacts ☁	14.6%	T1564.003: Hidden Window	4.4%
		T1564.001: Hidden Files and Directories	4.4%
		T1564.008: Email Hiding Rules ☁	3.9%
		T1564.011: Ignore Process Interrupts	1.7%
		T1564.012: File/Path Exclusions	1.7%

Maintain Presence

Defense Evasion (Continued)

T1218: System Binary Proxy Execution	11.2%	T1218.011: Rundll32	6.6%
		T1218.007: Msiexec	4.6%
		T1218.005: Mshta	0.5%
		T1218.002: Control Panel	0.2%
		T1218.014: MMC	0.2%
T1112: Modify Registry	10.0%		
T1036: Masquerading	9.2%	T1036.001: Invalid Code Signature	1.7%
		T1036.005: Match Legitimate Resource Name or Location	1.5%
		T1036.011: Overwrite Process Arguments	0.2%
		T1036.008: Masquerade File Type	0.2%
		T1036.003: Rename Legitimate Utilities	0.2%
T1553: Subvert Trust Controls	8.3%	T1553.002: Code Signing	8.3%
T1222: File and Directory Permissions Modification	7.5%	T1222.002: Linux and Mac File and Directory Permissions Modification	6.6%
		T1222.001: Windows File and Directory Permissions Modification	1.0%
T1574: Hijack Execution Flow	7.3%	T1574.011: Services Registry Permissions Weakness	6.8%
		T1574.001: DLL	0.5%
		T1574.008: Path Interception by Search Order Hijacking	0.2%
T1550: Use Alternate Authentication Material ☁	6.3%	T1550.001: Application Access Token	5.3%
		T1550.002: Pass the Hash	1.0%
T1202: Indirect Command Execution	4.6%		
T1666: Modify Cloud Resource Hierarchy ☁	4.6%		
T1006: Direct Volume Access	3.4%		
T1556: Modify Authentication Process ☁	2.7%	T1556.006: Multi-Factor Authentication ☁	1.5%
		T1556.003: Pluggable Authentication Modules	0.2%
		T1556.002: Password Filter DLL	0.2%
T1656: Impersonation ☁	1.7%		
T1578: Modify Cloud Compute Infrastructure ☁	1.5%	T1578.003: Delete Cloud Instance ☁	1.2%
		T1578.005: Modify Cloud Compute Configurations ☁	0.2%
T1055: Process Injection	1.0%		

Maintain Presence

Defense Evasion (Continued)

T1548: Abuse Elevation Control Mechanism ☁	1.0%	T1548.002: Bypass User Account Control	0.5%
		T1548.003: Sudo and Sudo Caching	0.2%
		T1548.001: Setuid and Setgid	0.2%
T1134: Access Token Manipulation	0.7%	T1134.001: Token Impersonation/Theft	0.7%
T1484: Domain or Tenant Policy Modification ☁	0.7%	T1484.001: Group Policy Modification	0.7%
T1207: Rogue Domain Controller	0.5%		
T1127: Trusted Developer Utilities Proxy Execution	0.2%	T1127.001: MSBuild	0.2%
T1220: XSL Script Processing	0.2%		
T1221: Template Injection	0.2%		
T1599: Network Boundary Bridging	0.2%	T1599.001: Network Address Translation Traversal	0.2%
T1610: Deploy Container	0.2%		
T1620: Reflective Code Loading	0.2%		
T1647: Plist File Modification	0.2%		
T1672: Email Spoofing ☁	0.2%		

Maintain Presence

Persistence

T1133: External Remote Services	23.8%		
T1098: Account Manipulation ☁	18.7%	T1098.007: Additional Local or Domain Groups	7.8%
		T1098.005: Device Registration ☁	3.6%
		T1098.001: Additional Cloud Credentials ☁	1.0%
		T1098.004: SSH Authorized Keys ☁	0.5%
		T1098.003: Additional Cloud Roles ☁	0.2%
T1078: Valid Accounts ☁	17.0%	T1078.004: Cloud Accounts ☁	9.5%
T1543: Create or Modify System Process	16.7%	T1543.003: Windows Service	8.3%
		T1543.002: Systemd Service	1.9%
		T1543.004: Launch Daemon	0.2%
T1505: Server Software Component	14.8%	T1505.003: Web Shell	14.8%
T1053: Scheduled Task/Job	11.2%	T1053.005: Scheduled Task	9.0%
		T1053.003: Cron	1.7%
T1136: Create Account ☁	10.2%	T1136.001: Local Account	3.2%
		T1136.003: Cloud Account ☁	1.5%
		T1136.002: Domain Account	1.0%
T1547: Boot or Logon Autostart Execution	8.3%	T1547.001: Registry Run Keys / Startup Folder	7.5%
		T1547.009: Shortcut Modification	1.0%
		T1547.013: XDG Autostart Entries	0.5%
		T1547.005: Security Support Provider	0.2%
T1574: Hijack Execution Flow	7.3%	T1574.011: Services Registry Permissions Weakness	6.8%
		T1574.001: DLL	0.5%
		T1574.008: Path Interception by Search Order Hijacking	0.2%
T1546: Event Triggered Execution ☁	6.1%	T1546.003: Windows Management Instrumentation Event Subscription	4.1%
		T1546.015: Component Object Model Hijacking	0.7%
		T1546.007: Netsh Helper DLL	0.5%
		T1546.010: AppInit DLLs	0.2%
		T1546.004: Unix Shell Configuration Modification	0.2%
		T1546.001: Change Default File Association	0.2%
T1671: Cloud Application Integration ☁	4.6%		
T1556: Modify Authentication Process	1.5%	T1556.006: Multi-Factor Authentication ☁	1.5%
T1037: Boot or Logon Initialization Scripts	0.5%	T1037.001: Logon Script (Windows)	0.2%
		T1037.004: RC Scripts	0.2%
T1176: Software Extensions	0.5%	T1176.001: Browser Extensions	0.2%
T1137: Office Application Startup ☁	0.2%	T1137.005: Outlook Rules ☁	0.2%
T1554: Compromise Host Software Binary	0.2%		

Maintain Presence

Command And Control

T1105: Ingress Tool Transfer	24.5%		
T1102: Web Service	15.5%	T1102.002: Bidirectional Communication	0.5%
T1071: Application Layer Protocol	11.2%	T1071.001: Web Protocols	3.4%
		T1071.004: DNS	3.2%
		T1071.002: File Transfer Protocols	0.7%
T1095: Non-Application Layer Protocol	10.9%		
T1572: Protocol Tunneling	8.0%		
T1573: Encrypted Channel	2.9%	T1573.002: Asymmetric Cryptography	2.9%
T1090: Proxy	2.7%	T1090.001: Internal Proxy	1.0%
		T1090.003: Multi-hop Proxy	0.5%
		T1090.002: External Proxy	0.5%
T1219: Remote Access Software	1.0%	T1219.002: Remote Desktop Software	0.5%
		T1219.003: Remote Access Hardware	0.5%
T1571: Non-Standard Port	0.7%		
T1659: Content Injection	0.2%		

Mission Completion

Collection

T1074: Data Staged ☁	39.6%	T1074.001: Local Data Staging	3.6%
		T1074.002: Remote Data Staging ☁	0.7%
T1213: Data from Information Repositories ☁	12.6%	T1213.002: Sharepoint ☁	6.1%
		T1213.004: Customer Relationship Management Software ☁	2.7%
		T1213.003: Code Repositories ☁	1.2%
		T1213.006: Data from Information Repositories: Databases ☁	0.5%
		T1213.005: Messaging Applications ☁	0.2%
T1560: Archive Collected Data	11.4%	T1560.001: Archive via Utility	5.6%
T1005: Data from Local System	11.2%		
T1114: Email Collection ☁	7.3%	T1114.003: Email Forwarding Rule ☁	0.7%
		T1114.002: Remote Email Collection ☁	0.2%
T1039: Data from Network Shared Drive	3.4%		
T1530: Data from Cloud Storage ☁	2.7%		
T1115: Clipboard Data	1.7%		
T1602: Data from Configuration Repository	1.2%	T1602.002: Network Device Configuration Dump	1.2%
		T1602.001: SNMP (MIB Dump)	0.2%
T1113: Screen Capture	0.5%		
T1025: Data from Removable Media	0.2%		
T1125: Video Capture	0.2%		

Mission Completion

Exfiltration

T1041: Exfiltration Over C2 Channel	7.8%		
T1567: Exfiltration Over Web Service ☁	5.8%	T1567.002: Exfiltration to Cloud Storage	2.9%
		T1567.001: Exfiltration to Code Repository	0.5%
		T1567.003: Exfiltration to Text Storage Sites	0.2%
T1020: Automated Exfiltration	0.5%	T1020.001: Traffic Duplication	0.2%
T1011: Exfiltration Over Other Network Medium	0.2%	T1011.001: Exfiltration Over Bluetooth	0.2%
T1030: Data Transfer Size Limits	0.2%		
T1052: Exfiltration Over Physical Medium	0.2%	T1052.001: Exfiltration over USB	0.2%

Impact

T1486: Data Encrypted for Impact ☁	17.7%		
T1489: Service Stop	10.7%		
T1657: Financial Theft ☁	10.4%		
T1565: Data Manipulation	6.1%	T1565.001: Stored Data Manipulation	3.9%
T1490: Inhibit System Recovery ☁	2.7%		
T1485: Data Destruction ☁	2.2%		
T1496: Resource Hijacking ☁	1.5%		
T1491: Defacement ☁	1.2%	T1491.002: External Defacement ☁	0.5%
		T1491.001: Internal Defacement	0.2%
T1529: System Shutdown/Reboot	1.2%		
T1531: Account Access Removal	1.2%		

Google Cloud

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