Mandiant

Google Cloud Security

M-Trends

2026 Report

Executive Edition

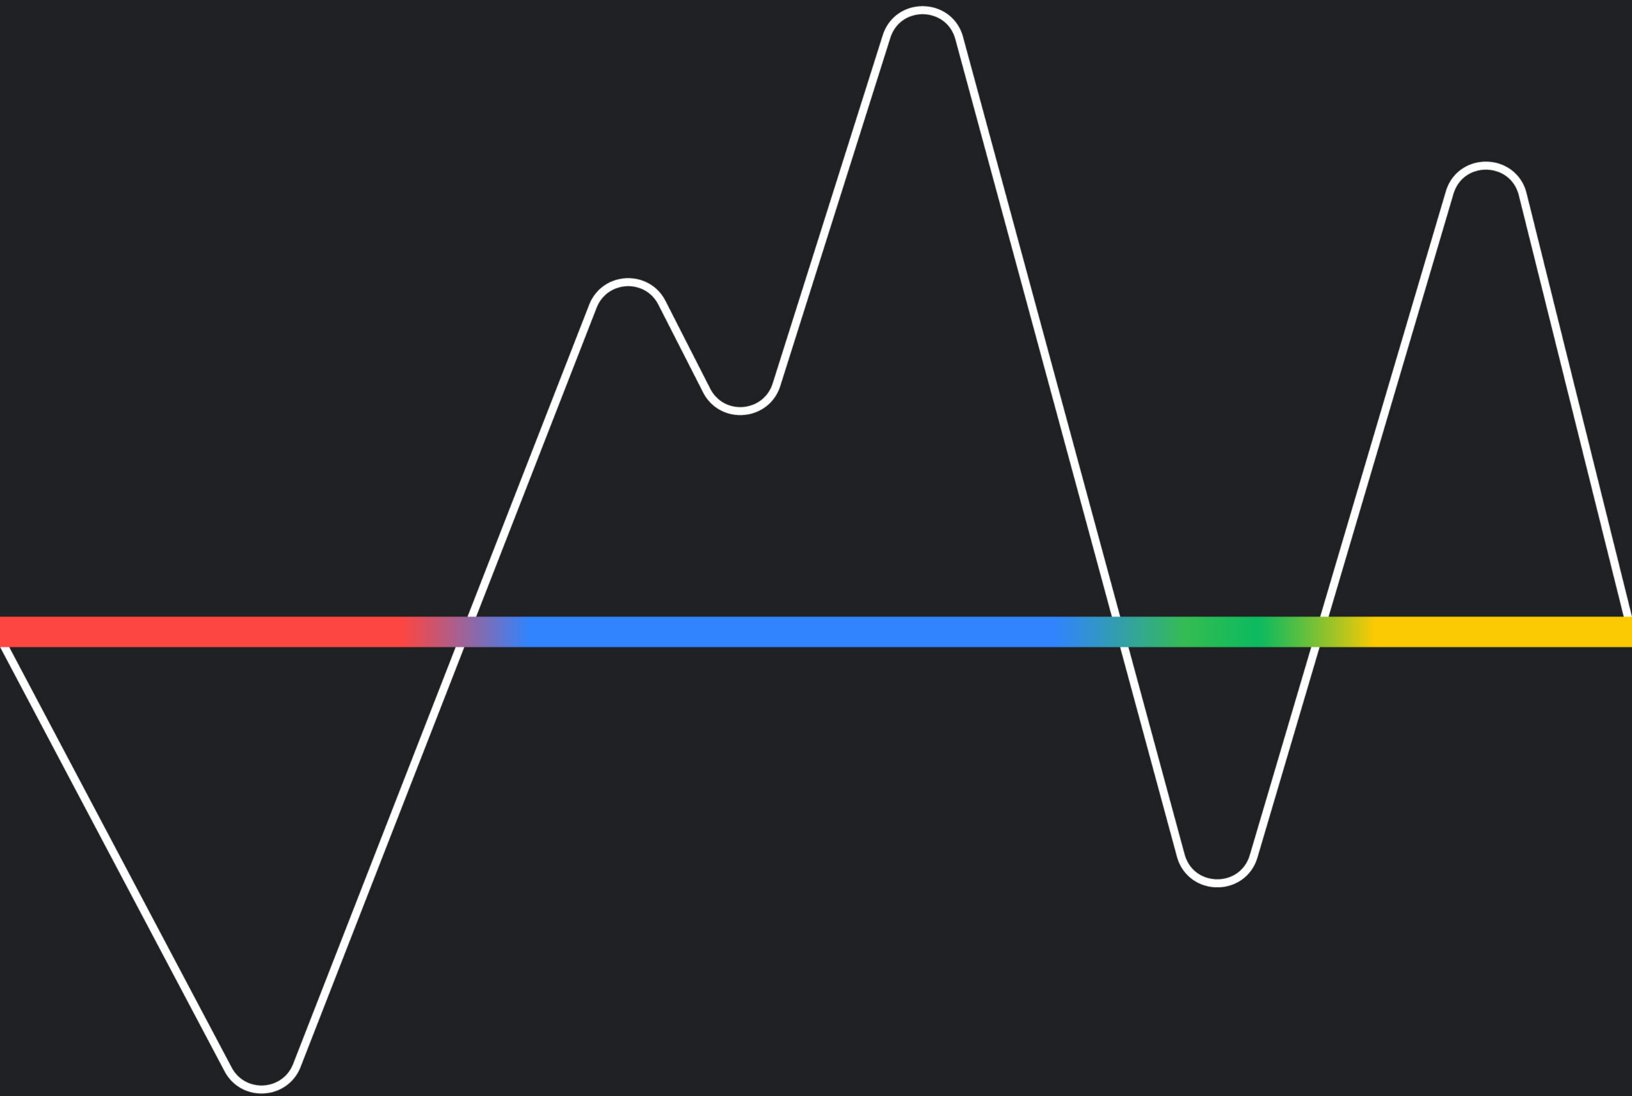

M-Trends 2026 Executive Edition

Foreword

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 are driving future attacks.

Recent GTIG reporting confirms adversaries are adopting AI. Threat actors are using large language models (LLMs) for hyper-personalized social engineering, malware that queries LLMs mid-execution to evade detection, and “distillation attacks” targeting proprietary machine learning logic. Mandiant red teams are incorporating AI-driven techniques into engagements to prepare organizations for these emerging threats; however, our M-Trends 2026 findings show that mitigating the human and systemic failures that enable breaches is mission critical.

A major takeaway from our 2025 incident engagements is that a subset of adversaries are remaining undetected on networks for longer periods of time, often by establishing persistence in edge devices that typically lack standard telemetry. Mandiant responded to enough of these types of incidents in 2025 that global median dwell time has risen to 14 days from 11 days in the previous reporting period, driven largely by long-term espionage and DPRK IT worker operations.

Other threat groups prioritize speed. We are observing a rising trend where initial access partners work directly with secondary groups rather than selling access on underground markets. This results in a “hand-off” that sometimes occurs in less than 30 seconds, creating a scenario where “minor” alerts can very quickly become major compromises.

Simultaneously, adversaries are systematically targeting infrastructure such as backups, identity services, and virtualization layers to deny recovery, putting immense pressure on organizations to pay ransom demands or risk losing the ability to recover.

To build true operational resilience, organizations must move at the speed of the adversary. A big part of that is understanding how adversaries are finding success. By closing critical visibility gaps and adopting defenses detailed in M-Trends 2026, enterprises can shift from reactive recovery to proactive containment before a minor alert becomes a catastrophic compromise.

By the Numbers

The metrics reported in M-Trends 2026 are based on Mandiant Consulting investigations of targeted attack activity conducted between Jan. 1, 2025 and Dec. 31, 2025.

The Bottom Line

Attackers are shifting their initial access strategies; exploits are still the most common, but rising to the number two spot is highly interactive, voice-based social engineering. At the same time, the global median dwell time has increased as sophisticated espionage groups and insider threats prioritize stealthy, long-term access. As adversaries look for opportunities to leverage and weaponize AI, exploit zero-day vulnerabilities on edge devices, and execute hand-offs between initial access partner and cybercrime groups, organizations should evolve beyond static defenses to continuously monitor identity behavior and infrastructure such as virtualization that has traditionally been outside the scope of EDR and other similar security tools.

What do I need to know?

- Financially motivated groups represented 41% of threat clusters observed in 2025 (a decline from 55% in 2024), while cyber espionage groups increased to 16% (up from 8%).
- The most common initial infection vector was exploits (32%) for the sixth consecutive year. Voice phishing saw a surge to 11% to become the second most commonly observed vector, while email phishing saw a steady decline from 14% in 2024 to 6% in 2025. While both email phishing and voice phishing fall under a broader social engineering umbrella, the distinction between the two is crucial for defenders: interactive attacks are significantly more resilient against automated technical controls and require different detection strategies.
- In ransomware-related incidents, prior compromise was the most frequently observed initial infection vector at 30%. Some threat clusters focus on gaining an initial foothold at many organizations via high volume, opportunistic infection vectors, then sell or hand off this access to other threat clusters for post-compromise exploitation.
- The most frequently targeted industries were high tech, financial, business and professional services, and healthcare, though the full scope of incidents affected more than 16 industry verticals. Notably, investigations in the high tech sector outpaced the financial sector organizations, which had the largest share in 2024 and 2023.
- Across all 2025 investigations, 52% of the time organizations first detected evidence of malicious activity internally (an increase from 43% in 2024), while 34% of the time they were notified by an external entity (down from 43%). Adversaries informed organizations of a compromise 14% of the time.
- In ransomware cases, adversaries notified the target 44% of the time, internal identification occurred 41% of the time, and external entities notified 15% of the time. A much higher percentage of adversary notifications is consistent with the ransomware business model.
- Global median dwell time rose to 14 days from 11 days in 2024. Comparing the distribution from 2024 to 2025 reveals a marginal decrease in incidents discovered in a week or less, and a shift toward longer dwell times (between one week and six months). This shift likely reflects the quantity of cyber espionage and North Korean IT worker incidents in which threat clusters prioritize maintaining stealthy, long-term access; both of these categories had a median dwell time of 122 days.
- Of the malware families observed in 2025 investigations, 36% were backdoors, 11% were downloaders, 10% were ransomware, 10% were droppers, and 9% were credential stealers.

- After five consecutive years as the most frequently observed malware family in Mandiant investigations, Cobalt Strike BEACON fell to fourth. The most frequently observed malware family was the GOLDVEIN.JAVA downloader, followed by REDBIKE (Akira) ransomware.
- Threat clusters are utilizing stealthy tactics and lightweight malware (such as the BRICKSTORM backdoor) on appliances that do not support EDR. Both financially motivated and cyber espionage groups abuse native functionalities in on-premises and cloud environments, as well as legitimate tools, to reduce opportunities for detection.
- Threat clusters have increasingly adopted AI tools to achieve productivity gains in reconnaissance, social engineering, and malware development. Furthermore, attackers are weaponizing AI within compromised environments; for example, the QUIETVAULT credential stealer was observed checking targeted machines for AI CLI tools to execute predefined prompts to search for configuration files.

What do we need to do?

- **Expand Visibility Beyond the Traditional Endpoint:** Deploy advanced threat detection across the entire ecosystem. Specifically, incorporate network traffic analysis for EDR-less edge appliances, and enforce strict telemetry for virtualization infrastructure.
- **Aggressively Manage Internet-Facing Attack Surfaces:** With exploits remaining the top initial infection vector for the sixth consecutive year, organizations should prioritize rapid patching, vulnerability scanning, and strict isolation of external-facing web application servers that are frequently targeted by zero-day and n-day campaigns.
- **Pivot Security Awareness Training Beyond the Inbox:** While email phishing remains a threat actor staple, Mandiant increasingly observes threat actors using other initial infection vectors, including interactive voice phishing, stolen credentials, and tactics like ClickFix. Educate employees and IT help desk staff specifically on recognizing live, voice-based social engineering, messaging app lures, and unauthorized MFA reset requests.
- **Shift to Continuous Identity Verification:** Because interactive social engineering frequently bypasses traditional MFA, and a variety of threat actors including North Korean IT workers seek to maintain long-term access, organizations should enforce strict least privilege, regularly audit SaaS/cloud integrations, and proactively hunt for anomalous identity behavior among both employees and remote contractors.
- **Update IR Playbooks for the Modern Extortion Pipeline:** Ensure incident response and recovery plans address both encryption-based ransomware and pure data-theft extortion. Conduct tabletop exercises focused on detecting opportunistic infections early to reduce response times and disrupt the hand-off from initial access partners to secondary groups.
- **Secure Developer Environments and AI Toolchains:** Threat clusters have been observed weaponizing AI within compromised environments. Because we have observed malware actively abusing legitimate local AI command-line tools to locate and steal GitHub and NPM tokens, organizations should adopt principles from the Google Secure AI Framework (SAIF). Specifically, by “extending detection and response” to include AI tools in the threat model, security teams can establish behavioral baselines and monitor for anomalous prompts or unauthorized data theft originating from these utilities.
- **Engage Red Teams for Modern Threat Emulation:** Regularly test defenses with realistic emulation of the latest adversary tactics. Measure the time it takes security teams to detect and respond to living-off-the-land techniques, edge device exploitation, and interactive social engineering.

A Minor Infection Today Can Be a Ransomware Attack Tomorrow

The Bottom Line

Closer collaboration between cybercriminal partners has collapsed the window for defense, shrinking the median time between opportunistic initial access from one group and the time at which a secondary threat group has access to just 22 seconds; down from previous years where it was closer to 8 hours. This shift mandates that organizations treat low-impact alerts as critical indicators, necessitating immediate remediation before high-impact actors can capitalize on the access.

What do I need to know?

- **The "Time to Hand-Off" Has Collapsed:** The median time between an initial access event and the hand-off to a secondary threat group dropped from more than 8 hours in 2022 to just 22 seconds in 2025. This decrease suggests a shift toward direct coordination where initial access partners establish access specifically for secondary groups to use immediately, rather than selling it later on underground channels.
- **Low-Impact Intrusions can be the Precursors to High-Impact Attacks:** Threat actors are using a division of labor model where specialized groups use low-impact techniques, like malicious advertisements (malvertising) or fake browser updates, to gain a foothold. Because these initial vectors look like low-impact malware, organizations hunting solely for high-impact tactics often miss them until it is too late.
- **Partner Pre-Staging:** Rather than simply selling an open door on underground markets, initial access partners now pre-stage the secondary group's preferred malware, tunnels, or backdoors during the initial infection. Bypassing the underground market and pre-configuring the environment allows follow-on actors to utilize these established footholds on their own timeline, meaning they are fully equipped to launch high-impact operations the moment they first interact with the network.

What do we need to do?

- **Treat Low-Impact Alerts as Critical Indicators:** The disparity between the perceived low criticality of initial infections and their potential high-impact outcomes requires a restructuring of response playbooks. Security teams should treat routine malware alerts as high-priority indicators of an impending hand-off to a secondary group.
- **Enforce a Baseline of Pre-Approved Tools:** IT and Security teams should define and deploy a set of pre-approved, centrally stored tools for users. This reduces the noise of opportunistic self-installs, and allows defenders to quickly identify and respond to deviations from the baseline.
- **Optimize for Event Correlation and Context:** Defensive workflows should shift focus from individual alert review to event correlation. Enriching alerts with contextual data allows analysts to spot the specific patterns of behavior that indicate a potential high-impact intrusion is underway.
- **Remediate Before Interactive Activity Begins:** The goal for defenders is to remediate an intrusion during the initial access phase—often a non-interactive event—before a secondary group begins hands-on-keyboard operations. Stopping the attack at this single-system stage is significantly more effective than attempting to recover from the subsequent high-impact activity.

Ransomware is Now a Resilience Problem

The Bottom Line

Ransomware groups are no longer just encrypting data; they are actively destroying the ability to recover. They target the most critical systems, effectively forcing a choice: pay or rebuild. True resilience now means designing networks so that recovery tools are segmented and protected. Locking down core systems and making it harder for attackers to move around inside the network gives security teams an advantage against today's ransomware threats.

What do I need to know?

- **Ransomware Has Evolved into Recovery Denial:** Ransomware operators have shifted their primary objective from simple data theft to recovery denial. Attackers now target system and administrative pathways—specifically identity services, virtualization management planes, and backup infrastructure—to reduce an organization's ability to recover while maximizing pressure to pay.
- **Identity Is the New Perimeter:** Sophisticated threat groups are manipulating the identity control plane to gain total control of targeted environments. We have observed attackers exploiting misconfigurations to issue certificates and create admin accounts that bypass multi-factor authentication (MFA) and password rotation. In some cases, attackers steal the entire databases or compromise cloud tenants to destroy infrastructure backends, locking defenders out of their own emergency accounts during a crisis.
- **Traditional Safety Nets Are Being Destroyed:** The strategy of relying on backups is failing as attackers actively hunt for and destroy backup architectures. Threat actors conduct reconnaissance to map storage locations and retrieve encryption configurations before systematically deleting backup objects from cloud storage and local systems. In on-premises scenarios, attackers have been observed unlinking virtualization environments from backup platforms or encrypting local recovery points, rendering standard incident response playbooks ineffective because the tools required to execute them are unavailable.

What do we need to do?

- **Increase Minimal Viable Security:** Organizations should treat identity as the primary perimeter. This requires hardening high-impact pathways by implementing Privileged Identity Management and Conditional Access to ensure administrative privileges are time-bound and observed. Security teams should enforce MFA on all access, utilize hardened Privileged Access Workstations, and audit Service Principal Names. Telemetry should be tuned to detect "Living Off the Land" behaviors to identify administrative takeovers before they impact the environment broadly.
- **Isolate Critical Tier-0 Control Planes:** Virtualization and management platforms should be treated as "Tier-0" assets with the strictest access constraints. Organizations should implement zero-trust segmentation to isolate these systems, and formally sever Active Directory (AD) integration to prevent a single identity compromise from translating into mass encryption. Utilizing dedicated out-of-band management with local, MFA-protected accounts ensures structural separation and prevents scenarios where production and recovery capabilities are lost simultaneously.
- **Advance Recovery Path Reliability:** Recovery capability should survive the compromise of the production environment. Organizations need to establish a dedicated, isolated recovery environment where teams can clean, validate, and stage system restores, helping mitigate a reinfection loop. This requires maintaining offline or immutable versions of Tier-0 assets—specifically identity and backup catalogs—that do not share the fate of the production network. Disaster Recovery plans should explicitly account for the total loss of the primary identity fabric.

Multi-Year Intrusions Highlighting Extreme Persistence

The Bottom Line

Prevention is ideal, but preparation is mandatory. Sophisticated threat actors are maintaining multi-year access by exploiting blind spots and administrative trust. If you cannot prove the scope of an intrusion due to logging gaps, you risk a loss of customer trust by being forced to assume and disclose a worst-case data theft. Treat visibility as a continuous audit to ensure you can detect and remediate these threats before they become unmanageable crises.

What do I need to know?

- **Pervasive Stealth and Persistence:** Certain threat actors are achieving dwell times exceeding one year, often by prioritizing legitimate credentials over custom malware. They target virtualization infrastructure and unmanaged edge devices, using built-in tools to blend in with standard administrative activity and evade detection.
- **Critical Visibility Gaps:** Current logging strategies fail to capture the full scope of these intrusions. With cases involving BRICKSTORM (a stealthy backdoor) averaging 393 days of dwell time, standard 90-day log retention leaves organizations unable to identify the initial access vector. Furthermore, reliance on EDR creates blind spots, as actors target network edge appliances and authentication protocols that sit outside the reach of endpoint agents.
- **Anti-Forensics Stifles Responders:** Sophisticated actors have been observed destroying the evidence needed to scope an incident. By employing techniques to mask file modifications and clearing system logs, threat actors make it nearly impossible to reconstruct the timeline of data theft. In the absence of definitive forensic proof, organizations risk severe regulatory and reputational damage by being forced to assume and disclose a worst-case-scenario breach.

What do we need to do?

- **Expand Log Retention and Scope:** Given the lengthy dwell times for certain sophisticated threat actors, standard 90-day retention policies are insufficient for scoping intrusions. Organizations should prioritize forwarding logs to centralized, long-term storage, specifically capturing data from blind spots like network edge devices, hypervisors, and authentication protocols that standard EDR tools often miss.
- **Secure Log Integrity:** To counter anti-forensics techniques like local log wiping and timestomping, prioritize forwarding logs to a centralized location immediately. This ensures evidence is preserved even if the host is compromised. Additionally, configure repositories to alert on unexpected log termination, which allows operations teams to identify when an attacker is intentionally disabling security controls to hide their tracks.
- **Shift to Proactive Hunting:** Security teams should move beyond reactive alerts and implement a routine of proactive threat hunting. This involves using advanced analysis techniques, such as stack ranking data, to identify outliers in authorized behavior. Integrating the latest threat intelligence allows teams to distinguish benign administrative activity from adversaries using native tools to hide in plain sight.
- **Validate Assets and Assumptions:** Leaders should enforce collaboration between security and infrastructure teams to audit environments. Regularly test assumptions to ensure critical assets are actually generating the expected telemetry. This process should also be used to identify and remove unused technology and systems to reduce the administrative burden and attack surface.

Adversary Focus on Virtualization Infrastructure

The Bottom Line

Virtualization platforms have shifted from backend infrastructure to frontline targets. Attackers are exploiting the “Tier-0” nature of hypervisors to bypass guest-level defenses, embed deep persistence that survives standard remediation, and deploy ransomware at a level that renders traditional recovery impossible. Protecting this stack requires treating the management plane as an isolated critical asset, and eliminating severe logging blind spots to restore visibility.

What do I need to know?

- **Hypervisors Are a Security Blind Spot:** Hypervisors often run proprietary operating systems that are incompatible with standard EDR tools. This creates a visibility gap where attackers can deploy malware or create unmanaged virtual machines to stage attacks without triggering EDR or SIEM alerts.
- **Attackers Are Bypassing Guest Defenses:** Sophisticated adversaries no longer need to log into target systems to steal data. By targeting the virtualization storage layer directly, they can clone virtual disks and extract sensitive Active Directory databases without ever interacting with the guest operating system or its security controls.
- **Ransomware Now Targets the Datastore:** Modern ransomware campaigns are moving down the stack, encrypting hypervisor datastores rather than individual machines. This tactic allows a threat actor to shut down and lock every server on a host simultaneously, rendering all associated virtual machines inoperable and unrecoverable.
- **Attackers Are Embedding Deep Persistence:** Adversaries are increasingly targeting the hypervisor’s underlying shell to deploy hidden backdoors and rogue virtual machines. Because these mechanisms run below the standard server environment, they routinely survive standard incident remediation efforts and system reboots, granting attackers permanent backdoor access.

What do we need to do?

- **Decouple Identity and Management:** Remove hypervisors and backup infrastructure from the corporate Active Directory domain. Relying on the same identity provider for both production and infrastructure creates a single point of failure. Deploy a dedicated Infrastructure-Only Identity Provider (IdP) or enforce Just-in-Time (JIT) access.
- **Isolate the Management Plane:** Treat virtualization interfaces as Tier-0 assets. Restrict management traffic to a dedicated, firewalled network segment accessible only through hardened Privileged Access Workstations, and enforce phishing-resistant MFA.
- **Enforce Immutable Resilience:** To counter the destruction of recovery capabilities, backup environments should be isolated and utilize immutable storage. Regular restoration testing should be conducted from these air-gapped copies to verify backups will function as intended during a high-pressure recovery event.
- **Centralize Infrastructure Telemetry:** Mandate the forwarding of virtualization management and hypervisor-level logs to the central SIEM. Because attackers are bypassing guest operating systems, hypervisor telemetry is now the only way to detect unauthorized access, snapshot manipulation, or rogue virtual machine creation.

Systematic Exploitation of Edge and Core Network Devices

The Bottom Line

Attackers are weaponizing edge and core network devices to evade modern security tools, exploiting vulnerabilities faster than patches are released and abusing native device features to silently steal data. Because these critical gateways are frequently uncatalogued and unmonitored, they grant adversaries invisible, long-term access. Organizations should urgently prioritize comprehensive asset discovery, strict patch management, and centralized logging to reclaim control of their network perimeter.

What do I need to know?

- **Zero-Day Exploitation is Accelerating:** The mean time to exploit vulnerabilities has plummeted, meaning threat actors are increasingly compromising systems, including edge and core network devices, before vendors even release a patch.
- **Network Devices are Security Blind Spots:** Because edge appliances cannot run traditional EDR software, attackers use them as safe havens. They deploy custom, in-memory malware that evades detection and survives in environments with poor telemetry.
- **Adversaries are Bypassing Endpoints:** Advanced attackers are conducting nearly the entire attack lifecycle—from reconnaissance to data theft—directly from network infrastructure, completely bypassing well-monitored workstations and servers.
- **Incident Response is Severely Hindered:** Edge devices have minimal storage and strict uptime requirements. When a compromise occurs, standard file system forensics is often impossible, making it difficult to confirm the presence of an attacker before evidence is lost.
- **Attackers Are Weaponizing Native Features:** Adversaries are increasingly using built-in administrative subshells and native packet-capturing functionality to collect copies of live traffic transiting the devices. This allows them to extract passwords from cleartext network protocols and operate entirely under the radar without needing to deploy detectable malware.

What do we need to do?

- **Centralize and Retain Network Logs:** Do not rely solely on endpoint data. Forward critical network device logs, especially application and administrative logs, to a centralized SIEM, retaining administrative logs for at least a year to ensure visibility during an investigation.
- **Enforce Strict Vulnerability Management:** Integrate network devices into comprehensive vulnerability scanning and asset management programs. Establish clear ownership and staggered patching schedules to remediate flaws before attackers can exploit them without disrupting business operations.
- **Develop Network-Specific Incident Playbooks:** Security teams should build detailed architectural diagrams and response playbooks. Pre-plan forensic collection procedures to ensure critical volatile evidence isn’t destroyed during a standard reboot or power-cycling event.
- **Identify Uncatalogued Network Infrastructure:** Because edge appliances are frequently deployed and forgotten, security teams should use comprehensive discovery scans to identify uncatalogued systems. Organizations cannot defend, patch, or monitor an asset they do not know exists.

The Cascading Impact of Third-Party SaaS Compromises

The Bottom Line

The shift to cloud-first infrastructure has transformed SaaS applications into pathways for massive supply chain attacks. Threat actors are bypassing standard defenses by stealing integration tokens and exploiting unvetted third-party apps, turning a single vendor breach into a cascading enterprise crisis. Organizations should shift to continuous identity verification, strictly governing end-user application consent, and enforcing rigorous third-party risk management before procurement.

What do I need to know?

- **MFA is Being Bypassed:** Attackers are no longer just stealing passwords; they are harvesting long-lived OAuth tokens and session cookies. Because these often remain valid post-logout, attackers can hijack sessions without triggering MFA alerts.
- **Integrations are the New Perimeter:** Threat actors compromise third-party SaaS vendors to steal hardcoded keys and personal access tokens, using those stolen secrets to seamlessly pivot into downstream customer environments and execute large-scale data theft.
- **Social Engineering is Escalating:** Financially motivated groups are aggressively targeting IT help desks with voice-based phishing (vishing). They impersonate employees to bypass controls, gain initial SaaS access, and mine the environment for highly privileged keys.

What do we need to do?

- **Discover and Govern the SaaS Estate:** Traditional asset management often misses cloud apps. Deploy SaaS Security Posture Management (SSPM) tools to actively inventory all applications, integrations, and hidden secrets (like API keys) to eliminate blind spots.
- **Harden Identity Controls:** Mandate that all SaaS applications route through a central Identity Provider (IdP). Enforce strict least privilege for third-party API keys, eliminate wildcard permissions, and utilize Just-in-Time (JIT) access to minimize standing privileges.
- **Automate Lifecycle Management:** Implement automated rotation for all secrets and service account credentials. Enforce extremely short lifespans for access tokens and browser sessions to instantly devalue stolen cookies sold on the dark web.
- **Lock Down End-User App Consent:** Administrators should disable the ability for end users to consent to unverified third-party applications. This ensures that only vetted applications can obtain persistent tokens, effectively blocking the path for malicious OAuth applications to gain access to the environment.
- **Enforce Strict Vendor Risk Management:** A risk accepted in a third-party application is a risk to the production core. Organizations should implement a robust Third-Party Risk Management program to vet vendors before onboarding, mandating features like single sign-on, granular audit logging, and secure development practices as part of the procurement process.

Download the full report.

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

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