



Global Incident Response Report 2026

Executive Summary

WE SEE FOUR MAJOR TRENDS THAT WILL SHAPE THE THREAT LANDSCAPE FOR 2026.

First, AI has become a force multiplier for threat actors. It compresses the attack lifecycle, from access to impact, while introducing new vectors. This speed shift is measurable: in 2025, exfiltration speeds for the fastest attacks quadrupled.

Second, identity has become the most reliable path to attacker success. Identity weaknesses played a material role in almost 90% of Unit 42 investigations. Attackers increasingly log in with stolen credentials and tokens, exploiting fragmented identity estates to escalate privileges and move laterally.


Third, software supply chain risk has expanded beyond vulnerable code to the misuse of trusted connectivity. Attackers exploit software-as-a-service (SaaS) integrations, vendor tools and application dependencies to bypass perimeters at scale. This shifts the impact from isolated compromise to widespread operational disruption.

Fourth, nation-state actors are adapting stealth and persistence tactics to modern enterprise operating environments. These actors increasingly relied on persona-driven infiltration (fake employment, synthetic identities) and deeper compromise of core infrastructure and virtualization platforms, with early signs of AI-enabled tradecraft used to reinforce these footholds.

While these four trends each present a challenge, **attacker success is rarely determined by a single attack vector.** In more than 750 incident response (IR) engagements, 87% of intrusions involved activity across multiple attack surfaces. This means defenders must protect endpoints, networks, cloud infrastructure, SaaS applications and identity together. Further, nearly half (48%) involved browser-based activity, reflecting how often attacks intersect with routine workflows like email, web access and day-to-day SaaS usage.

Most breaches were enabled by exposure, not attacker sophistication. In fact, in over 90% of breaches, preventable gaps materially enabled the intrusion: limited visibility, inconsistently applied controls, or excessive identity trust. These conditions delayed detection, created paths for lateral movement, and increased impact once attackers obtained access.

Security leaders must close the gaps attackers rely on. First, reduce exposure by securing the application ecosystem, including third-party dependencies and integrations, and hardening the browser, where many intrusions now begin. In parallel, reduce area of impact by advancing zero trust and tightening identity and access management (IAM) to remove excessive trust and limit lateral movement. Finally, as the last line of defense, ensure the security operations center (SOC) can detect and contain threats at machine speed by consolidating telemetry and automating response.






Table of Contents

INTRODUCTION	04
EMERGING THREATS AND TRENDS	06
INSIDE THE INTRUSION	20
RECOMMENDATIONS FOR DEFENDERS	28
APPENDIX	36
METHODOLOGY	42



SECTION 1: Introduction

In 2025, Unit 42 responded to more than 750 major cyber incidents. Our teams worked with large organizations facing extortion, network intrusions, data theft and advanced persistent threats. Targets spanned every major industry and more than 50 countries. In each case, the situation had escalated to the point where the SOC called for backup.

When that call comes, our incident responders move quickly to investigate, contain and eradicate the threat. We help organizations establish what happened, restore operations, and reduce the risk of recurrence by strengthening controls, visibility and resilience.

Each intrusion tells a story: what the attacker targeted, how they gained access, how the activity escalated and what could have stopped it sooner. In the aggregate, these stories become trends and provide insight into the global threat landscape. They show what's changing in adversary tradecraft, the repeated mistakes organizations make, and most importantly, what defenders can do to keep their organizations safe. This report distills those lessons.

Over the past year, attack speeds continued to accelerate. Attackers are still early in their adoption of AI-enabled tradecraft, but its impact is already visible. AI reduces friction across reconnaissance, social engineering, scripting, troubleshooting and extortion operations. It enables greater scale and the ability to launch multiple attacks simultaneously. The result is a shrinking window for detection and containment, where what happens in the first minutes after initial access can determine whether an incident becomes a breach.

At the same time, most breaches still follow familiar paths. And that is why our most important conclusion remains unchanged: security is solvable. In more than 90% of incidents, misconfigurations or lapses in security coverage materially enabled the intrusion. Attackers are adapting, but they most often succeed by exploiting preventable gaps—inconsistent control deployment, incomplete telemetry, over-permissive identity trust and unmanaged third-party connectivity across SaaS and cloud.

This report is organized as a practical guide to the current threat landscape:

Emerging Threats and Trends:

How attacker tradecraft is evolving—AI as a force multiplier, identity as the most reliable path to success, expanding software supply chain risk through trusted connectivity and evolving nation-state tactics.

Inside the Intrusion:

An aggregate view of observed tactics, techniques and procedures across Unit 42 investigations—what attackers target, how they get in, how fast they move and the impacts they drive.

Recommendations for Defenders:

Concrete steps to close the gaps that enable compromise, constrain area of impact, and build response capability fast enough to stop incidents before they escalate.

Unit 42 operates 24/7 to protect the digital world from cyberthreats. The goal of this report is straightforward: to turn what we learn on the front lines into decisions that stop incidents before they become breaches.




Sam Rubin
SVP of Consulting and Threat Intelligence
Unit 42

02

SECTION 2:

Emerging Threats and Trends



TREND 1:

AI Has Become a Force Multiplier for Attackers

AI is changing the economics of intrusions. It increases attacker speed, scale and effectiveness while opening entirely new attack vectors.

While much of this activity occurs on adversary infrastructure—beyond our ability to directly observe—Unit 42 investigations and research reveal a clear shift. In 2025, threat actors moved from experimentation to routine operational use. AI is not an attacker “easy button,” but it is a massive friction reducer. It allows threat actors to move faster, iterate more frequently, and operate with fewer human constraints.

AI INCREASES THE SPEED AND SCALE OF ATTACKS

AI compresses the attack lifecycle and reduces the manual effort required to operate across multiple targets.

Faster vulnerability exploitation: The window between disclosure and exploitation continues to shrink. Threat actors are automating the “monitor → diff → test → weaponize” loop. Unit 42 research found that attackers start scanning for newly discovered vulnerabilities within 15 minutes of a CVE being announced. Exploitation attempts often begin before many security teams have even finished reading the vulnerability advisory.

Parallelized targeting: Operator time is less of a constraint. AI-assisted workflows allow actors to run reconnaissance and initial access attempts across hundreds of targets in parallel, and then concentrate effort where they find a weak signal.

Ransomware at scale: We see actors using AI to reduce manual work during deployment (script generation, templating) and extortion (messaging consistency). The shift is not that ransomware is new, it is that the operator time required to run it at scale is dropping.

What this means in time-to-impact: Last year, Unit 42 simulated an AI-assisted attack that reduced time-to-exfiltration down to 25 minutes. Real-world IR data reflects this acceleration: The fastest 25% of intrusions reached exfiltration in 72 minutes, down from 285 minutes the calendar year prior.

Ransomware Automation

In a ransomware investigation, Unit 42 recovered operational scripts used to deploy payloads, coordinate lateral movement and impair security controls at scale. Several elements were consistent with AI-assisted development, including unusually thorough commenting, templated variants and efficiency-focused fallback logic. The net effect was machine-like execution across hundreds of systems, compressing the time and effort typically required to stage a multi-phase deployment.

The AI Negotiator

In an extortion case, Unit 42 negotiators observed responses that were unusually consistent in tone, grammar, cadence and turnaround time across exchanges. These patterns are consistent with templated or AI-assisted messaging. Even partial automation matters: it enables actors to run more concurrent negotiations and apply more disciplined pressure, without tying up a human operator on every thread.

AI IMPROVES ATTACKER OUTCOMES

AI is raising the success rate of known attack techniques.

Hyper-personalized social engineering: We have moved past “phishing with better grammar.” Actors can automate open-source intelligence (OSINT) collection, including professional and organizational context, to craft lures that match the target’s role and relationships.

Synthetic identities: Threat actors like [Muddled Libra](#) and [North Korean](#) IT workers increasingly use deepfake techniques to steal credentials and pass remote hiring workflows.

Malware development: In the [Shai-Hulud](#) campaign, Unit 42 assessed that attackers used a large language model (LLM) to generate malicious scripts.

Lowered barrier to entry: Purpose-built malicious LLMs and jailbreak attacks continue to reduce the skill required to produce persuasive lures and functional code variants. The net effect is that more actors are able to execute credible tradecraft faster, with fewer mistakes.

Vibe Extortion

An unsophisticated actor exfiltrated sensitive data but had no plan for the shakedown. To bridge the gap, they used an LLM to script a professional extortion strategy, complete with deadlines and pressure tactics. The result was surreal: The actor recorded a threat video from their bed while visibly intoxicated, reading the AI-generated script word-for-word from a screen. The threat lacked technical depth, but the model supplied coherence. AI didn’t make the attacker smarter; it just made them look professional enough to be dangerous.

Bottom line: AI improves the attackers’ rates of success at each stage. It improves the quality of lures, shortens the time needed to adapt tools and reduces dependence on constant operator intervention, making extortion more consistent and scalable.



AI CREATES NEW ATTACK VECTORS

Enterprise AI adoption creates a new class of risk: Living off the AI land (LOTAIL). Just as attackers misuse PowerShell or Windows Management Instrumentation (WMI), they are now weaponizing legitimate AI platforms and embedded assistants.

Turning your AI platform into a weapon: Threat actors use valid credentials to misuse enterprise AI platforms. For example, recent Unit 42 [research](#) on Google Vertex AI demonstrated how attackers could misuse custom job permissions to escalate privileges and use a malicious model as a Trojan horse to exfiltrate proprietary data.

The attacker’s co-pilot: With compromised credentials, an intruder can use an internal assistant to pull context at machine speed, including requesting integration guides, admin runbooks or network maps. The assistant becomes a force multiplier, allowing intruders to understand the environment with fewer mistakes.

The AI-Assisted Insider

An insider weaponized their company’s own AI assistant to stage an attack. Forensic analysis showed the insider used the tool to research internal systems, generate a custom denial-of-service (DoS) script and troubleshoot errors in real time. The assistant bridged a skill gap, enabling the actor to target core infrastructure they likely could not have operated against as effectively without AI support.

The risk is clear: If a tool can help employees get work done, it can also help intruders understand your environment and move with fewer mistakes.



COUNTERMEASURES: DEFENDING AGAINST AI-DRIVEN THREATS

These tactics will help you
defend against AI-assisted attacks

Counter AI-accelerated attack speed

Automate external patching:

Mandate automated patching for critical CVEs on internet-facing assets to close the 24-hour exploitation window.

Autonomous containment:

Deploy AI-driven response to drive down mean time to detect/respond (MTTD/MTTR) and isolate threats before they can automate lateral movement.

Defend against improved tradecraft

Behavioral email security:

Transition from signature-based filters to engines that identify anomalies in communication patterns.

Intent-based awareness:

Move beyond simply training employees to spot typos. Shift to out-of-band (OOB) verification for all sensitive requests (e.g., wire transfers, credential resets or remote hiring).

Protect the AI attack surface

Monitor model telemetry:

Correlate unusual AI API calls or scripts sourced from model outputs with known evasion techniques.

Prompt visibility:

Alert on sensitive queries to internal LLMs (e.g., “find all passwords”) and enforce strict permission boundaries for tokens and service accounts.

TREND 2:

Identity Is the Most Reliable Path to Attacker Success

In the past year, identity weaknesses played a material role in nearly all (90%) of the investigations Unit 42 handled. In our caseload, identity shaped intrusions end to end. It served as the way in, the path to privilege escalation and the mechanism for lateral movement using valid access.

As organizations move deeper into SaaS, cloud and hybrid environments, the network perimeter matters less. Identity—the linkage between users, machines, services and data—has become the practical perimeter. In many cases, threat actors don't need a sophisticated exploit chain. They log in with stolen credentials, hijacked sessions or mis-scoped privileges.

Authenticated access changes the dynamics of an intrusion. It lets adversaries move faster, blend into normal activity and expand their area of impact with fewer obstacles. This trend is accelerating as machine identities, embedded AI applications and fragmented identity estates expand the number of access paths attackers can exploit.

THE WAY IN: IDENTITY-DRIVEN INITIAL ACCESS

Unit 42 case data shows that 65% of initial access is driven by identity-based techniques, as shown in Figure 1. While defenders focus on patching vulnerabilities, threat actors often bypass software controls by targeting users and authentication paths.




Figure 1. The data highlights identity as the dominant driver of initial access in modern intrusions.

We see the following primary routes to initial access:

Identity-related social engineering (33%): Identity-based phishing (22%) and other social engineering (11%) remain the leading drivers of modern breaches. Rather than simple credential theft, these tactics increasingly focus on multi-factor authentication (MFA) circumvention and session hijacking, allowing attackers to bypass authentication controls and move laterally by exploiting trusted identity workflows.

Credential misuse and brute force (21%): Previously compromised credentials (13%) and brute force activity (8%) allow attackers to gain access with little interaction. By using valid accounts obtained from prior breaches or underground markets, actors log directly into virtual private networks (VPNs), remote access gateways and cloud portals, bypassing traditional perimeter defenses without triggering early detection.

Identity policy and insider risk (11%): Stemming from internal trust and architectural flaws, these vectors involve the exploitation of valid permissions. Attackers leverage IAM misconfigurations (3%), such as overly permissive policies, to escalate privileges and inherit access, while insider threats (8%) involve the abuse of legitimate credentials.

Identity and vulnerability management are not separate fights. A leaked credential can create the same exposure as an unpatched internet-facing system.

THE WAY THROUGH: IDENTITY TURNS ACCESS INTO IMPACT

After initial access, identity gaps are one of the most common ways attackers turn a foothold into a high-impact breach. In modern environments, authenticated actions determine speed and blast radius.

Unit 42 analysis of more than 680,000 identities across cloud accounts found that 99% of cloud users, roles and services had excessive permissions, some unused for 60 days or more. This creates an environment where lateral movement is easier than it should be, because many identities carry privileges they don't need day to day.

Attackers exploit both human and machine identities as operational levers:

Privilege escalation: Over-scoped roles, inherited permissions and unretired legacy grants create repeatable paths to higher privilege. Once an attacker can write to IAM, they can often escalate quickly without deploying novel tooling.

Credential reuse and lateral movement: Actors commonly test compromised credentials across other systems. This is especially true where passwords are reused across production and non-production environments, or where shared accounts still exist.

Token and OAuth misuse: Stolen session tokens and illicit OAuth grants let attackers bypass interactive authentication (including MFA), persist without repeated logins and operate with fewer obvious alerts.

Trust paths (e.g., shared administrative accounts, delegated access and third-party tools) become fast lanes for lateral movement. Without tight privilege boundaries and strong identity segmentation, a single compromised identity can expand into broad access.

99%
of 680,000
cloud identities
had excessive
permissions



THE EXPANDING IDENTITY ATTACK SURFACE

The identity landscape is expanding and fragmenting. As organizations adopt cloud, SaaS and AI-enabled workflows, identity moves into areas that often sit outside consistent governance, creating areas where attackers operate with reduced visibility.

Three trends are driving this shift:

The rise of machine and AI identities: Non-human identities, like service accounts, automation roles, API keys and emerging AI agents, often outnumber human users. These identities are frequently over-privileged, rely on long-lived credentials and are inconsistently monitored. For an attacker, compromising a service account can be higher leverage and quieter than compromising a person.

Shadow identities: Cloud and AI adoption has increased the volume of unsanctioned accounts, developer environments and third-party connectors. These shadow identities often bypass standard onboarding, review and logging, creating access paths the SOC might not see until after impact.

Identity silos: Most enterprises operate multiple identity systems (e.g., Active Directory, Okta, cloud-native IAM). When authentication and authorization are fragmented, so is visibility. Attackers can move between on-premises and cloud environments while leaving incomplete trails in any single control plane.

Misconfiguration at scale turns identity from a control into a liability. When machine identities, shadow access and fragmented identity estates combine, attackers gain more reliable paths to persist and expand. And defenders lose end-to-end visibility.

COUNTERMEASURES: DISRUPTING IDENTITY-DRIVEN TRADECRAFT

These tactical steps can disrupt the identity-related tradecraft observed in Unit 42 cases.

Deploy phishing-resistant MFA: Standard MFA is not enough against modern bypass and adversary-in-the-middle tactics. Prioritize FIDO2/WebAuthn hardware keys or passkeys for high-value roles (admins, executives, developers).

Inventory and rotate machine identities: Establish continuous discovery for non-human identities (service accounts, automation roles, API keys). Immediately rotate static credentials for any privileged service account that has not changed in 90 days and reduce credential lifetime wherever possible.

Harden the session: Attackers increasingly pivot post-login by stealing tokens and misusing OAuth grants. Reduce session lifetimes for sensitive applications and enforce conditional access that continuously evaluates device health, location and risk during the session.

Eliminate standing admin rights: Move privileged access to a just-in-time model. Remove persistent admin grants and require time-bound elevation with approvals and strong logging, so a compromised account yields minimal privilege by default.

TREND 3.

Software Supply Chain Attacks Increasingly Drive Downstream Disruption

Supply chain risk is no longer limited to vulnerable code. In 2025, the supply chain expanded to include SaaS integrations, vendor management planes and complex dependency ecosystems. The defining pattern was downstream disruption and parallel assessment. When an upstream provider reported a compromise or outage, customers were often left to stop and answer a basic question: are we affected? In many cases, they had limited visibility into their own exposure.

The new failure mode is not one compromised customer. There are many customers pushed into parallel triage while the upstream picture is still unclear. This makes the supply chain a high-value target for both nation-states and criminal groups. A single compromise can create a one-to-many opportunity, delivered through the trusted connectivity modern business relies on.

SAAS INTEGRATIONS: INHERITED PERMISSIONS AT SCALE

SaaS environments are stitched together through OAuth apps, API keys and workflow automation. These connections routinely carry access to data and business processes. For attackers, compromised integrations can become a lateral movement path that looks like normal automation.

This exposure is reflected in Unit 42 investigations. Data from SaaS applications was relevant to 23% of cases in 2025, up from 18% in 2024, 12% in 2023, and just 6% in 2022. The steady increase shows how attackers are moving past traditional perimeters and concentrating on the cloud-based tools where modern work now takes place.

The risk is inherited permissions. When an organization integrates a third-party app via OAuth, that application receives whatever rights were originally granted, sometimes including the ability to read sensitive data, manage users or modify records. If the upstream provider is compromised, those same permissions can be misused downstream.

The Hidden Integration Risk

In a recent investigation involving a compromised sales engagement platform (Salesloft/Drift integration), attackers leveraged valid OAuth tokens to access downstream Salesforce environments. The activity resembled routine customer relationship manager (CRM) automation and blended into expected integration traffic. Post-incident review revealed a deeper issue: the organization discovered nearly 100 additional third-party integrations connected to Salesforce, many dormant, unmonitored or owned by former employees.



OPEN SOURCE AND AI: DEPENDENCY SPRAWL AND BUILD-TIME COMPROMISE

Open source remains the foundation of modern development, but the risk increasingly concentrates in indirect dependencies. [Unit 42 research](#) indicates that over 60% of vulnerabilities in cloud-native applications reside in transitive libraries, which are the “silent” dependencies pulled in through packages that your code relies on.

Threat actors are also injecting malicious code into upstream packages to execute during install and build steps, compromising pipelines before deployment. Development velocity compounds this risk. As GenAI-assisted coding becomes mainstream, teams are ingesting more code and more dependencies faster. This is often done with insufficient scrutiny of provenance, maintainer trust and downstream package behavior.

The Malicious Package

We investigated a campaign where threat actors uploaded malicious versions of legitimate npm packages. One package, embedded deep in a dependency tree, executed attacker-controlled code immediately upon installation. Because this activity occurs during build and install, it can bypass runtime detections and establish a foothold across multiple build environments before anyone sees an alert.



VENDOR TOOLS: WEAPONIZING MANAGEMENT CHANNELS

Vendor tools, especially remote monitoring and management (RMM) and mobile device management (MDM) platforms, are designed for privileged administrative action at scale. When attackers gain access to a vendor's management infrastructure (or the customer's tenant), they can push malware, run commands or change configurations in ways that blend into routine administrative traffic. This trend is backed by our observations in the field: We identified that 39% of command-and-control (C2) techniques were related to remote access tools (T1219).

Enterprises also inherit risk from opaque third-party applications running inside critical workflows. When customers cannot inspect a vendor's codebase or security assumptions, latent backdoors, hard-coded credentials or exposed interfaces can persist unnoticed.

The Legacy Application

In a multi-national investigation, a legacy third-party billing application exposed an undocumented, unauthenticated interface to the internet. Existing controls did not detect it because the traffic appeared consistent with normal application behavior. Deeper assessment revealed structural flaws, including SQL injection points and hidden shell functionality. These issues had persisted for years because the customer could not inspect the underlying code.



THE IMPACT: FROM RESPONSE TO BUSINESS DISRUPTION

Supply chain incidents amplify disruption through uncertainty. When a supplier is compromised, downstream teams operate in an information vacuum. The result is organizations going into “assessment mode” at scale, as teams pause changes, review integrations, isolate dependencies and attempt to confirm the absence of impact before normal operations resume.

Three systemic gaps drive that burden:

Inventory gaps: Many organizations lack a unified view of SaaS connections, vendor agents and transitive libraries, slowing the answer to “where is this used?”

Permission opacity: Effective privileges of integrations, agents and tooling are difficult to determine quickly without manual review, making the true impact unclear.

Telemetry gaps: Because activity arrives through trusted channels (updates, API calls, administrative tooling), logs often look legitimate, which can delay detection and increase investigation time.

Looking ahead: This challenge will compound as organizations adopt AI-enabled workflows and third-party agents. Supply chain risk will increasingly include not just code integrity, but the integrity of models, connectors and delegated actions executed on an organization’s behalf.

COUNTERMEASURES: SAFEGUARDING THE SOFTWARE SUPPLY CHAIN

Defending the supply chain requires reducing the time needed to assess exposure and the area of impact.

Map SaaS ownership and scope: Inventory OAuth apps and integrations (SaaS security posture management and discovery). Assign owners. Remove dormant integrations and those tied to departed users.

Design “break-glass” severing plans: Predefine how to revoke tokens, disable connectors and isolate vendor agents without improvising during an upstream incident.

Log vendor and integration activity at audit depth: Ensure you can answer what was executed, where and by whom. Alert on permission changes, token grants and anomalous admin actions.

Harden build ingestion: Use software composition analysis (SCA) and provenance controls. Pin versions, restrict new repositories and require review for new dependencies, especially those that execute at install or build time.

TREND 4.


Nation-State Actors Are Adapting Tactics to Modern Environments

Nation-state operations expanded in 2025, advancing espionage, pre-positioning and access campaigns. Across campaigns affiliated with China, North Korea and Iran, three shifts stood out:

- Greater use of identity-driven access
- Deeper compromise of infrastructure and virtualization layers
- Early experiments with AI-enabled tradecraft aimed at stealth and persistence.

China-aligned groups moved beyond user-level activity into infrastructure and virtualization platforms. North Korean and Iranian operators broadened their use of recruitment lures, synthetic personas and tailored malware to establish access. We also observed emerging AI-driven techniques, including deepfake identity creation and automated C2 generation.

These developments reflect a shift toward access methods that are significantly harder for defenders to detect and validate.



CHINA: FOCUSED ON THE EDGE AND VIRTUALIZATION

Chinese-nexus threat activity continued to prioritize long-term access and data collection. Notable shifts in 2025 moved from email-focused espionage to deeper exploitation of application, infrastructure and virtualization layers.

Phantom Taurus exemplified this change, evolving from campaigns centered on sensitive email collection to direct targeting of databases and web servers for collection and exfiltration. Its NET-STAR malware used advanced evasion techniques, posing significant risk to organizations with exposed web infrastructure.

Similarly, we observed a year-long persistence campaign against information technology, SaaS and business-process outsourcing organizations (tracked by Unit 42 as activity cluster CL-STA-0242). The group behind the campaign compromised virtualization platforms operated by IT service providers and deployed BRICKSTORM malware, which concealed C2 traffic inside ordinary encrypted web sessions, making detection through network monitoring far more difficult. CISA has publicly attributed BRICKSTORM activity to China state-sponsored actors.

These shifts illustrate a continued move away from user-level collection toward deeper compromises of infrastructure and virtualized environments, where long-term access is both more durable and harder for defenders to detect.

NORTH KOREA: WEAPONIZED HR PART I

North Korean threat activity remained a persistent challenge for enterprises in 2025. Multiple long-running campaigns continued despite extensive public reporting, law-enforcement actions and multilateral sanctions measures.

Unit 42 tracked at least two campaigns:

Wagemole: North Korean operatives obtained unauthorized remote employment with U.S. and European organizations and covertly routed income back to the regime. The access gained through these contractor and employee roles enabled both unauthorized financial payments and espionage. First publicly exposed in 2023, Wagemole remained active in 2025, and we identified and evicted related activity from more than 20 enterprise environments.

Contagious Interview: Since at least 2022, operators have targeted software developers and IT personnel through fictitious job interviews that deliver malware via coding challenges. In 2025 alone, we removed Contagious Interview infections from more than 10 enterprise networks, underscoring the risks associated with running unverified code on corporate systems.

IRAN: WEAPONIZED HR PART II

Iranian threat activity remained high in 2025 as multiple groups continued operations against strategic sectors. Of particular note are Screening Serpens and Curious Serpens, both of which used employment-themed lures to target aerospace and satellite-communications providers. This activity reflects Iran's long-running interest in organizations that handle sensitive technical and operational information.

Unit 42 tracked the following campaigns:

Screening Serpens (aka Smoke Sandstorm, UNC1549): This group targeted government organizations in the Middle East by creating fraudulent employment portals that mimicked well-known aerospace and defense companies. These sites delivered malware packaged as job application materials, often signed with valid code-signing certificates to increase credibility. Operational security errors allowed Unit 42 researchers to review the full infection chain, which prompted candidates to download an infected survey file or document bundle.

Curious Serpens (aka APT33, Peach Sandstorm): Curious Serpens targeted a communications provider through job-recruitment lures sent by email and posted on career-oriented websites. The operation installed a modular backdoor capable of collecting intelligence and staging follow-on payloads. Operators relied on legitimate signed executables, DLL side-loading and evasion techniques, showing continued investment in a specialized tool set designed to circumvent modern security controls.

Using realistic employment themes and signed binaries increases the likelihood that victims will open malicious files. This highlights the need for sensitive sectors to monitor recruitment-related activity and verify any externally sourced documents or code.

The Trusted Résumé

In one Screening Serpens investigation, an attacker approached an employee through LinkedIn and personal email with a tailored résumé file that installed malware and enabled using legitimate remote-management tools. Once inside, the operator gathered credentials, surveyed the environment, deployed a custom backdoor and attempted to remove activity traces, indicating an emphasis on persistence and stealth.

NATION-STATE ADOPTION OF ARTIFICIAL INTELLIGENCE

Evidence of large-scale AI adoption by nation-state actors remains limited, but 2025 offered early signs that some groups are beginning to integrate AI into their operations. Much of this activity is difficult for defenders to observe, since many likely use cases (such as malware development, infrastructure generation or analysis of exfiltrated data) occur outside enterprise environments and beyond conventional visibility. As capabilities advance, understanding where nation-states are experimenting with AI is increasingly important for anticipating future tradecraft.

Attackers appear most interested in using AI to strengthen persistence and build more durable footholds. Nation-state operators have shown a growing reliance on identity and credibility-driven entry points and deeper compromise of virtualized and application infrastructure. These access methods are already difficult for defenders to validate, and AI will likely make them more efficient and harder to disrupt.

One of the clearest public examples emerged in July, when Ukrainian authorities, in a [CERT-UA advisory](#), reported that suspected Russian malware known as LAMEHUG used an LLM to generate C2 instructions through an API. Attributed to Fighting Ursa (aka APT28, Fancy Bear), the activity replaced a human operator with an automated workflow.

North Korean operators also showed signs of AI experimentation. In Unit 42 [research](#) associated with the Wagemole campaign, investigators identified suspected North Korean accounts using AI-based image manipulation services to create deepfake personas for employment fraud schemes. In a related Contagious Interview-style operation, attackers fabricated an entire company and populated it across multiple social networking platforms using AI-generated identities, repurposed accounts and modified profiles belonging to real professionals. The result was a convincing corporate façade designed to increase trust and improve the success rate of recruitment-driven access operations.

COUNTERMEASURES: DEFENDING AGAINST NATION-STATE ADVERSARIES

Focus defenses on the access paths, infrastructure layers and trusted channels nation-state operators use to gain and maintain long-term access.

Tighten verification across identity and recruitment workflows:

Strengthen checks on contractor onboarding and external hiring to catch synthetic personas, deepfakes and job-themed lures before they reach core systems.

Expand monitoring across virtualized and application infrastructure:

Baseline and log activity on virtualization platforms, web-facing applications and service-provider environments. Alert on deviations that signal persistence or lateral movement.

Harden and monitor the use of trusted tools and channels:

Review how signed binaries, encrypted traffic, remote-management tools and collaboration platforms are used. Flag patterns that suggest credential misuse or covert activity.

Instrument and govern AI-related activity in sensitive workflows:

Limit which AI services can interact with identities, source code or sensitive data. Log their use and investigate anomalous patterns that could indicate automated persona creation or AI-driven operations.

SECTION 3:

Inside the Intrusion

This section breaks down the behavior we observed in Unit 42 Incident Response investigations in 2025. We organize these observations into four dimensions to show what attackers are doing and how they are succeeding:

The Attack Surface:

This is where attackers strike. Intrusions rarely stay in one lane; they now span endpoints, cloud infrastructure and identity layers simultaneously.

The Entry Point:

This is how they get in. Phishing and vulnerabilities have tied as the leading initial access vectors, each at 22%. Attackers are pragmatic, they exploit human error and unpatched systems with equal frequency to force the door open.

The Velocity:

This is how fast they move. While average times vary, the fastest group of attackers is accelerating, shrinking the window for effective defense.

The Impacts:

This is the cost to the victim. This year marked a shift away from encryption and toward data theft and extortion.



3.1.

The Attack Surface: Intrusions Span the Enterprise

ATTACKS RARELY STAY IN ONE LANE

Figure 2 lists the primary attack surfaces involved in Unit 42 investigations in 2025, spanning endpoints, networks, cloud services, identity systems, applications, email and user-driven activity. These categories represent the primary operational layers where we observed attacker activity during investigations. Because intrusions frequently span multiple layers, they are not mutually exclusive and do not sum to 100%. A single incident may involve several at once.

Across all incidents, 87% involved activity across two or more attack surfaces. Sixty-seven percent of incidents involved activity across three or more surfaces. Activity across four or more attack surfaces appeared in 43% of attacks, and we have observed cases with activity across as many as eight attack surfaces. While the distribution of affected attack surfaces varies year to year, this pattern reinforces the fact that intrusions rarely remain confined to a single surface and often expand as access and opportunity grow.

Identity featured prominently in many incidents—at nearly 90%—representing one of the most commonly involved attack surfaces in our caseload.

Activity targeting humans also appeared frequently, accounting for 45% of incidents. This pattern echoes the broader themes in our recent [Social Engineering Report](#), which highlights how human-layer interaction continues to play a decisive role in intrusion success.




Figure 2. Attack surfaces involved in intrusions, showing the percentage of incidents in which each surface was affected.

THE BROWSER ATTACK SURFACE: ATTACKS AT THE HUMAN INTERFACE

Browser activity played a role in 48% of investigations this year (up from 44% in 2024). This reflects how routine web sessions expose users to malicious links, credential-harvesting pages and injected content when local controls are weak.

In one ClickFix incident we investigated, attackers directed an employee at a global industrial firm to a spoofed website through search engine optimization (SEO) poisoning while searching for a restaurant. The site used social-engineering prompts to convince the employee to execute malicious code copied into their clipboard, after which the attacker attempted to run malware in memory. The attacker appeared to be trying to download an infostealer, although we could not confirm the exact payload.

A global medical technology firm experienced an intrusion that began with SEO poisoning. An administrator accessed a spoofed site hosting a malicious version of an administrative tool, and the link was later shared with a domain administrator through an internal messaging call. This resulted in the execution of the compromised software. After gaining a foothold, the attacker deployed ransomware across key systems, exfiltrated data and issued a ransom demand. The resulting disruption affected manufacturing, distribution, shipping and order processing for an extended period while systems were restored.

Unmanaged applications and limited browser protections allowed an initial execution attempt in one incident before it was contained. In another, privileged execution of a malicious administrative tool enabled ransomware deployment and broader operational disruption.



THE CLOUD ATTACK SURFACE: COMPROMISING THE PIPELINE

Reflecting a continuation of last year's pattern, about 35% of our investigations involved cloud or SaaS assets. In these cases, the investigation required collecting logs or images from cloud environments or reviewing activity within externally hosted applications, indicating that the intrusion touched cloud-hosted assets or workflows.

Cloud weaknesses varied, but even basic issues shaped attacker behavior once they established access. In one investigation, sensitive cloud credentials were found exposed in a public repository, expanding the paths attackers could use to reach cloud environments.

In another investigation, attackers targeted a developer in an open-source forum and persuaded them to download a poisoned debugging tool. This turned a routine collaboration into a point of cloud compromise.








The compromised tool provided attackers with access to the developer's stored cloud credentials. They used these credentials to reach backend systems and trigger unauthorized withdrawals across several blockchain networks. This case shows how access obtained through cloud-native development workflows can be misused to reach sensitive systems and cause substantial impact.


3.2.

The Entry Point: Initial Access Comes from Predictable Paths

Initial access in 2025 followed a familiar pattern, with most intrusions beginning through a concentrated set of well-understood vectors. Figure 3 shows the distribution of those pathways across the past five years, highlighting how phishing and software vulnerabilities consistently appear among the top entry points. While the relative balance between vectors shifts year to year, the overall trend is stable: attackers continue to rely on a small number of dependable techniques to gain their initial foothold.

Figure 3: Initial access vectors (2021–2025). Unit 42 data collection methodology has adjusted to provide more granularity, reducing the “Other” category. Increased granularity also introduces new categories. When data is not available for a specific year, it is denoted by N/A.





PHISHING AND VULNERABILITIES TIE FOR DOMINANCE

Phishing and vulnerability exploitation are the most common initial access vectors, with each accounting for 22% of the initial access across 2025 incidents. This parity exists simply because both methods work incredibly well.

Phishing campaigns are achieving higher conversion rates as AI helps attackers craft credible, error-free lures that bypass traditional filters and engage users more effectively. At the same time, vulnerability exploitation is accelerating as attack surfaces expand and automation allows adversaries to scan for and exploit weaknesses faster than defenders can patch. Because both vectors offer a reliable path to compromise, attackers are heavily utilizing both.

Beyond phishing and vulnerability exploitation, we see important trends for the other key initial access vectors across the five-year dataset:

- Previously compromised credentials declined to 13% in 2025, reversing heightened activity reported in 2023 and 2024.
- Activity within the “Other Social Engineering” category grew substantially over the period, rising from 3% in 2021 to 11% in 2025 even after we introduced more granularity. Much of this growth appears to align with direct-interaction tactics such as the help-desk manipulation techniques used by groups like Muddled Libra.
- Brute force fell from 13% to 8%, ending a multi-year rise and suggesting stronger identity controls across many organizations.
- IAM misconfigurations remain a persistent initial access vector, appearing between 1% and 4% throughout the five-year period.

VULNERABILITY EXPLOITATION IS DRIVEN BY OPPORTUNITY, NOT NOVELTY

Attackers rely on vulnerability exploitation when it offers a clear operational advantage. The five-year pattern shows actors responding directly to the kinds of weaknesses available to them and the effort required to turn those weaknesses into access.

When high-impact issues appear in widely deployed systems, operators move quickly because the potential reach is substantial and the work needed to automate exploitation is relatively low.

This pattern reflects attacker pragmatism. Operators tend to exploit whatever is most accessible and cost-effective at any given moment.

BIG ENVIRONMENTS, BIGGER VULNERABILITY EXPOSURE

The data suggests that the largest enterprises face a different balance of initial-access risk: in 2025, vulnerabilities accounted for just over a quarter (26%) of initial access in these environments, compared with 17% for phishing. This pattern indicates that larger firms may be reducing their phishing exposure through stronger email filtering, user awareness and identity controls. These measures do not eliminate phishing risk but likely limit its effectiveness relative to smaller organizations.

Large, distributed environments with mixed ownership, legacy systems and uneven patching cycles make it easier for exploitable weaknesses to persist even in well-funded organizations. For firms of this size, complexity itself increases the likelihood that vulnerabilities go unaddressed, explaining why exploitation appears more frequently as an initial access vector.

3.3.

Velocity: The Fastest Attacks Are Getting Faster

The time-to-exfiltration, which measures the duration between initial compromise and confirmed data theft, shows a sharp acceleration at the fastest end of the spectrum. The quickest quartile of intrusions reached exfiltration in just over an hour (72 minutes) in calendar year 2025, down from nearly five hours (285 minutes) in 2024, as shown in Figure 4. The share of incidents reaching exfiltration in under one hour also increased—from 19% in 2024 to 22% in 2025.

Across the full dataset, the median time to exfiltration (MTTE) was two days. Although longer than the fastest incidents, even the median highlights how quickly attackers can access and remove data once inside the environment.

Defenders must be prepared for intrusions that progress from compromise to exfiltration in minutes or hours as well as slower, more methodical operations that unfold over days that involve deeper reconnaissance and durable persistence.

In 2024 it took
285 minutes
to exfiltrate

In 2025 it took
72 minutes
to exfiltrate

Figure 4. First-quartile attack speeds increased when comparing calendar year 2024 with calendar year 2025.

3.4.

The Impact: Extortion Beyond Encryption

Encryption appeared in 78% of extortion cases in 2025, a sharp decline from the near-or-above-90% levels for 2021–2024 shown in Figure 5. This represents the most pronounced year-over-year change in the dataset and shows that traditional ransomware has not disappeared, but it is no longer uniformly present in extortion operations.

The reduction in encryption does not correspond to a rise in other individual tactics. Instead, it reflects that attackers increasingly view encryption as optional rather than essential. Several 2025 intrusions proceeded with extortion even when victims retained access to their systems. In these cases, data exposure, direct pressure or both were sufficient to generate leverage without file-locking.

Data theft remained a consistent feature of extortion activity, appearing in more than half of cases year over year. Threat actors frequently used the threat of exposure on leak sites, and in some instances the resale of stolen data, to pressure victims regardless of whether encryption occurred.

Harassment, while less common, remained a persistent tactic. These behaviors included contacting employees directly, threatening to publish internal information or claiming they would sell customer data to other actors if victims didn't pay. Some groups escalated pressure by reaching out to customers or partners, amplifying reputational and operational strain even when systems remained accessible.

These patterns show that extortion has decoupled from encryption. While encryption remains prominent, attackers now have multiple reliable ways to create leverage. This broadens the range of conditions under which extortion can occur. It also reinforces the need for visibility, rapid response and strong data-handling practices regardless of whether attackers deploy ransomware.





Figure 5. How extortion tactics have changed from 2021–2025.

DATA THEFT REMAINS DURABLE LEVERAGE

Ransom economics helps explain why attackers continue to pursue these operations. Figure 6 shows that median initial demands increased from \$1.25 million in 2024 to \$1.5 million in 2025, and median payments also rose.

Median Initial Ransom Demands



Median Ransom Payments




Figure 6. Ransomware remains a lucrative option for attackers.

When measured against perceived annual revenue (PAR), these demands represented 0.55% of PAR, down from 2% the prior year. Many ransomware groups appear to be researching victims' ability to pay and using this information to calibrate demands. Asking for a lower percentage of PAR could reflect a strategy aimed at increasing the likelihood of payment.

Among organizations that chose to pay, median payments rose from \$267,500 to \$500,000, though payments as a share of PAR fell from 0.6% to 0.26%. The gap between initial demands and final payments shows how much room victims often have to negotiate, and it underscores the value of structured negotiation in limiting financial exposure.

The choice to pay remains highly situational, influenced by operational impact, regulatory considerations, legal requirements and business continuity needs. In 2025 cases where negotiations occurred, the median reduction between initial demand and final payment increased from 53% to 61%. This demonstrates how frequently experienced negotiators can reduce costs even as overall attacker pricing trends upward.

Many ransomware groups now operate with business-like structures including defined roles, affiliate programs and repeatable negotiation playbooks. Some cultivate "brand reputation" through dark web communications, portraying themselves as predictable or professional counterparts.

This brand maintenance extends to promise-keeping: in our 2025 dataset, threat actors fulfilled their commitments (such as providing decryption keys or allegedly deleting stolen data) in 68% of cases where they made a promise. For defenders, these recognizable patterns can provide leverage, though they never eliminate the risk of engaging with criminal actors.

Recovery practices also shape extortion outcomes. About 41% of victims were capable of restoring systems from backup without needing to pay, which reduced the operational impact of encryption but did not eliminate downtime. Even with recovery, many organizations still faced system rebuilds, containment work and other delays before returning to normal operations. Restoration is also fragile: in 26% of extortion cases, attackers impacted backups, adding further disruption.

When encryption is mitigated through backup restoration, or when backups fail entirely, the threat of exposure continues to pressure victims, ensuring data theft remains central to extortion activity.

04

SECTION 4.

Recommendations for Defenders

This section identifies the systemic weaknesses that enable attacks and the practical steps required to stop them. By addressing the root causes rather than just their symptoms, organizations can elevate their defenses to withstand both common and emerging threats.

4.1.

Common Contributing Factors: Why Attacks Succeed

Attacker success is rarely about zero-day exploits. Across the incidents we responded to in 2025, we found that **in more than 90% of incidents, preventable gaps in coverage and inconsistently applied controls directly contributed to the intrusion.**


These gaps determine how easily an attacker gains initial access, how quickly they move laterally, and whether defenders can detect and respond in time. Across this year's investigations, three systemic conditions appeared repeatedly.

1.

VISIBILITY GAPS: MISSING CONTEXT DELAYS DETECTION

Many organizations fail to leverage the telemetry needed to observe early-stage attacker behavior. Critical indicators of initial access and early attacker activity often go unnoticed because the SOC has not operationalized signals across endpoint, network, cloud and SaaS layers. The result is missing context: defenders might see individual events, but lack the correlation to recognize an active intrusion.

This fragmentation forces responders to manually reconstruct attacks from disparate tools, creating delays that attackers exploit. **In 87% of incidents, Unit 42 investigators reviewed evidence from two or more distinct sources to establish what happened, with complex cases drawing on as many as 10.** A lack of unified visibility consistently slowed detection, allowing adversaries to begin lateral movement before defenders could see the full picture.



2.

ENVIRONMENTAL COMPLEXITY: INCONSISTENCY CREATES THE PATH OF LEAST RESISTANCE

Security baselines are rarely applied universally. Over time, environmental drift, driven by legacy systems, technology adoption or merger and acquisition activity, makes it difficult to enforce a consistent standard across the enterprise.

In multiple investigations, critical controls like endpoint protection were fully deployed in one business unit yet missing or degraded in another. This inconsistency creates a path of least resistance. Over 90% of data breaches were enabled by misconfigurations or gaps in security coverage, rather than novel exploits.

3.

IDENTITY: EXCESSIVE TRUST LEADS TO LATERAL MOVEMENT

Across our investigations, identity weaknesses repeatedly turned an initial foothold into broader access. The core issue was often excessive trust—privileges and access paths that were too permissive or remained in place long after they were needed.

Attackers escalated privileges by misusing unretired legacy roles and over-permissioned service accounts. Rather than breaking in, they advanced by using valid access where the organization had left too much trust behind.

These failures reflect identity drift. As permissions accumulate and exceptions persist, intruders encounter fewer barriers. Nearly 90% of incidents trace back to an identity-related element as a critical source of the investigation or a primary attack vector.

4.2.

Recommendations for Defenders

The recommendations that follow focus on practical steps to address the systemic conditions described above.



1.

EMPOWER SECURITY OPERATIONS TO DETECT AND RESPOND FASTER

With the fastest attacks now exfiltrating data in roughly an hour, security operations must move at machine speed. This comes from empowering the SOC with comprehensive visibility across the enterprise, AI to identify the signal in the noise, and automation to drive immediate response and remediation. Adopting these six capabilities will put your SOC in the best position to succeed:

Ingest all relevant security data. Attackers do not operate in silos, yet defenders often monitor in them. In 2025, visibility gaps—particularly across SaaS, cloud identity and automation layers—were a primary driver of attacker success. Critical telemetry often existed but remained trapped in disparate systems, preventing defenders from correlating identity shifts with automation outputs or browser-stored artifacts like session tokens.

To detect modern intrusions, organizations must ingest and normalize signals from identity providers, cloud platforms and SaaS applications into a unified view. This consolidation closes the weak spots attackers exploit, allowing defenders to identify escalation routes early. Whether using rule-based detection or AI, the quality of insight depends entirely on the completeness of the data feeding it.

Prevent, detect and prioritize threats with AI-driven capabilities. High alert volumes and fragmented tools allow attackers to hide by spreading activity across systems. Without correlation, these actions appear unrelated, delaying escalation. AI-driven capabilities are essential to stitch these disparate signals into a unified operational view. Behavioral analytics help surface subtle anomalies, such as unusual token use or lateral movement through cloud automation, that rule-based detection often fails to catch.

AI strengthens defense by correlating events across identity, endpoint, cloud and network layers, prioritizing high-fidelity incidents over background noise. This allows security teams to distinguish coordinated attacks from routine activity instantly, ensuring analysts focus their efforts on the threats that pose the greatest risk rather than chasing false positives.

Enable real-time threat response with automation. Delays in containment often stem from unclear ownership and manual validation steps that cannot keep pace with attacker automation. Effective response requires assigning explicit authority for automated containment actions, such as revoking tokens or isolating workloads, so that execution can proceed without hesitation.

By replacing ad hoc judgment with standardized, validated playbooks, organizations ensure that response follows an auditable sequence. However, to meet the pace of modern threats, agentic AI must be deployed as the ultimate defense accelerator. These autonomous systems dynamically investigate complex alerts, correlating data across domains at machine speed to gain a complete picture.

Once validated, agents are authorized to execute dynamic, surgical containment actions, from isolating affected systems via microsegmentation to automatically revoking compromised credentials. This disciplined, intelligent approach dramatically reduces operational drift, limits attacker dwell time and prevents isolated compromises from escalating into broader incidents.

Transition from reactive to proactive security. To shift from reactive defense, organizations must move beyond traditional pentesting to continuous adversarial testing. Point-in-time audits rarely capture the interplay of identity drift and cloud misconfigurations that attackers exploit in real-world intrusions. Defenders need to validate how controls perform under realistic conditions, ensuring telemetry pipelines and response workflows operate as intended.

Proactivity extends to recovery. Resilient organizations verify that systems are free of residual access, such as compromised credentials or altered configurations, before restoring services. Ensuring that remediation addresses root causes, rather than simply restoring outdated snapshots, helps prevent rapid reinfection and supports long-term resilience.

Uplevel the SOC for high-performance outcomes. During active incidents, inconsistent containment or unclear ownership creates openings for attackers to re-establish access. High-performance SOC teams eliminate this variance by ensuring response actions are applied uniformly, regardless of the analyst or time of day. Consistency under pressure is critical; it prevents isolated compromises from escalating into broader crises.

Achieving this requires bridging operational silos across Security, IT, and DevOps. Playbooks should reflect how systems operate today, rather than how they were originally designed, so that automated actions align with real business logic. Empowering analysts with broader responsibility, such as end-to-end incident response rather than alert triage alone, improves retention, increases versatility and drives measurable business outcomes.

Deepen your bench with an IR retainer. The right retainer extends your capabilities beyond emergency response. To stay ahead, organizations must test and validate controls against the specific behaviors threat actors use in the wild. Recurring assessments across offensive security, AI security, SOC processes and cloud security help confirm that telemetry pipelines and response workflows operate as intended under realistic attack conditions.

Your IR retainer partner should provide rapid access to specialists for proactive readiness checks, detection engineering and validation, ensuring that defensive improvements hold up over time. By pairing continuous testing with retained expertise, organizations improve resilience.

By aligning your SOC with these core principles, you transform your defense into a high-velocity response engine capable of outmaneuvering adversaries and stopping threats before they escalate.

2.

ADOPT ZERO TRUST TO CONSTRAIN THE AREA OF IMPACT

Zero trust is a strategic necessity in an environment where identity has become the primary attack surface. The goal is to eliminate implicit trust relationships between users, devices and applications and to continuously validate every stage of a digital interaction.

In reality, achieving zero trust is complex. However, even small gains will reduce the attack surface, constrain lateral movement and minimize the impact of any initial access to your environment. By removing the assumption of safety inside the perimeter, defenders force attackers to work harder for every inch of access, slowing their velocity and creating more opportunities for detection.

Continuously verify users, devices and applications. Attackers frequently exploit the static trust that persists after an initial login. Once inside, they use stolen session tokens or valid credentials to masquerade as legitimate users, often bypassing perimeter controls entirely. Static checkpoints at the front door are no longer sufficient.

Continuous verification treats trust as dynamic, with decisions revisited as conditions change during a session. Validating identity context, device health and application behavior in real time allows organizations to detect when a legitimate session is hijacked or when user behavior deviates from the norm. As a result, compromised accounts or devices remain useful to attackers for only a limited period, reducing opportunities to expand access or stage data.

Enforce least privilege to constrain attacker movement. Excessive permissions act as a force multiplier for attackers. In many 2025 incidents, intruders bypassed internal controls by taking advantage of identity drift, using accumulated privileges and unretired roles that organizations failed to remove. Rather than relying on complex exploits, they moved laterally through valid but over-provisioned access paths.

Enforcing least privilege reduces this attack surface by limiting users, services and applications to only the access required for their function. This must extend beyond human users to include machine identities and service accounts, which often retain broad, poorly monitored permissions. Removing unnecessary rights eliminates the straightforward access paths attackers rely on, forcing them into more visible and difficult techniques that are easier for defenders to detect.

Apply consistent inspection across trusted and untrusted traffic. Apply consistent inspection across trusted and untrusted traffic. Attackers know that while the perimeter is guarded, internal “east-west” traffic between workloads often passes without inspection. They exploit this trust by using encrypted internal connections to move laterally and stage data without triggering alarms.

To achieve consistent, pervasive threat analysis, organizations must consolidate all network, cloud and secure access service edge (SASE) security onto a single unified platform. This unified fabric delivers consistent Layer 7 inspection everywhere, automatically enforcing policy via one management plane.

This consolidation enables the strategic shift to advanced cloud-delivered security services. This shift allows real-time, inline analysis of all traffic, including crucial decryption and inspection of traffic moving between internal workloads. This capability removes the spots where attackers hide, proactively stopping unknown phishing, zero-day malware and evasive C2 activity.

Control data access and movement to reduce impact. The most damaging outcomes in many incidents occur not at initial compromise but during subsequent data access, staging and exfiltration. Attackers often search for repositories with weak controls or poorly monitored flows to quietly aggregate sensitive information before detection.

Stronger governance over how data is accessed, shared and transferred reduces these opportunities by limiting where sensitive information can move and under what conditions. When data pathways are tightly controlled and consistently monitored, attackers face fewer options to prepare or extract valuable assets, reducing the scale and severity of potential loss even when a compromise occurs.

By systematically eliminating implicit trust, you strip attackers of the mobility they rely on, ensuring that a single point of compromise leads to a contained incident rather than an enterprise-wide crisis.

3.

STOP IDENTITY ATTACKS WITH STRONGER IDENTITY AND ACCESS MANAGEMENT

Identity is now the security perimeter, yet it too often remains poorly secured. Identity weaknesses were a determining factor in over half of the intrusions investigated in 2025, primarily because identity stores expanded faster than the controls intended to govern them.

Attackers consistently moved through the gaps created by this governance drift, exploiting legacy permissions and unmonitored service accounts to bypass perimeter defenses. To stop this, organizations must manage identity not as a static list of credentials, but as a dynamic operational asset across the entire lifecycle.

Centralize identity management for humans and machines. You cannot govern what you cannot see. When identity data is fragmented across legacy directories; cloud providers and SaaS environments, attackers take advantage of the resulting weak spots.

Centralizing user and machine identities into authoritative directories simplifies authentication and removes hidden access paths that are difficult to monitor consistently. This consolidation should also include third-party integrations and API connectors so that every entity requesting access, whether a person, a service account or an AI agent, is visible to security teams. With a unified control plane in place, defensive AI can correlate login anomalies with suspicious activity, turning identity into an active operational signal rather than a static list of credentials.

Combat governance drift with continuous lifecycle management. Governance drift, where operational changes move faster than the controls designed to guide them, remained a significant contributor to attacker leverage.

Role transitions, rapid deployment cycles and everyday shortcuts widened the gap between written policy and actual access. Permissions held by workflow tools and service connectors often exceeded what policy intended. This created escalation paths that attackers exploited through legacy permissions and unmonitored service accounts. Treating identity as a lifecycle, by limiting automation to current needs and retiring excess access over time, helps close these gaps and restrict attacker movement after initial access.

Detect and respond to identity-based threats. Defensive AI performs most effectively in environments where identities are managed as operational assets rather than static credentials. In our investigations, organizations with strong identity foundations showed earlier linkage between login anomalies, automation activity and peripheral identity events, which contributed to faster containment.

Where governance was strong, detection pipelines produced clearer and more reliable indicators that helped teams identify escalation behavior earlier. In contrast, weak governance created noise that obscured these signals. Regular reviews keep permissions aligned to real requirements, improving the accuracy of detection signals and ensuring that AI-assisted controls operate effectively.

Secure AI and automation integrity. As organizations embed AI agents and automated workflows into core processes, these systems become attractive targets for manipulation. In our investigations, we observed assistant accounts deployed with broad default access and automation tools running without integrity validation.

To prevent these tools from becoming vectors for attack, security teams must apply the same governance rigor to AI systems as they do to human users. This includes explicitly validating automation steps before they enter production, applying integrity checks to AI-enabled workflows and ensuring that assistant accounts are hardened against misuse.

By treating identity as a dynamic operational system rather than a static directory, you eliminate the hidden pathways attackers rely on and enable security teams to detect misuse the moment it occurs.

SECURE THE APPLICATION LIFECYCLE FROM CODE TO CLOUD

Protecting the modern enterprise requires more than securing infrastructure. It requires securing the factory that builds it.

In 2025, attackers increasingly targeted the software supply chain and cloud APIs to bypass traditional perimeters, injecting vulnerabilities into code or exploiting weak integrations before they ever reached production. To counter this, organizations must extend security safeguards from the earliest stages of development through to runtime, treating AI models, build pipelines and third-party code with the same rigor as internal systems.

Prevent security issues from reaching production. Security must operate at the speed of development. Integrating safeguards into DevOps and continuous integration and continuous deployment (CI/CD) pipelines helps identify and remediate vulnerabilities in custom code, open-source components, and AI configurations before deployment.

The same approach applies to AI systems, where early assessment of model security and configuration reduces downstream risk. Hardening development tools and governing open-source dependencies helps eliminate weak spots that attackers exploit to inherit trust within business workflows.

Secure the software and AI supply chain. Although not the most common attack vector, supply chain compromises yield the highest impact, especially for otherwise mature organizations. Weaknesses in build systems, integration services and AI-related repositories allow attackers to reach downstream environments without ever interacting with a firewall.

Reducing this exposure requires strict provenance checks. Build environments and deployment pipelines must have clear identity controls and integrity protections. External software libraries, API connectors and AI components should be evaluated for access patterns and update practices before adoption. Effective supply chain governance gives detection processes a reliable baseline, making it easier to identify when a trusted dependency begins behaving unexpectedly.

Identify and block runtime attacks. Once applications are live, the focus shifts to containment. Attackers frequently attempt to persist and expand access by misusing legitimate cloud identities, APIs or workload permissions.

Real-time detection, combined with consistent runtime controls such as behavioral monitoring, clear network boundaries and limits on unexpected API interactions, helps disrupt these tactics. The same protections should extend to AI hosting environments, where monitoring for model drift and unauthorized data access limits attacker movement even after initial compromise.

Automate cloud detection and response. In the cloud, speed is the only metric that matters. Delays in isolating affected workloads or revoking misused identities give attackers the room they need to escalate.

Automation allows SecOps teams to detect and respond to cloud-based threats continuously, using native cloud controls to contain incidents quickly. Actions such as isolating compromised containers or revoking suspicious session tokens help prevent localized issues from escalating into broader outages or data loss.

Build a culture of secure AI and development. AI is now an operational asset, not just a tool. As assistants and automated prompts become embedded in daily workflows, they introduce behavioral risks that technical controls alone cannot solve.

A strong security culture treats AI systems with the same discipline as critical infrastructure. This includes reviewing how assistants are used, avoiding the exposure of sensitive data in prompts and validating AI-generated code. When teams understand that human judgment remains central to effective AI use, governance controls are reinforced rather than bypassed, ensuring that the drive for automation does not outpace the ability to oversee it.

By embedding security into the fabric of your development and runtime environments, you help ensure that the speed of AI and cloud innovation drives business growth rather than systemic risk.

5.

SECURE THE ATTACK SURFACE AND THE HUMAN INTERFACE

Securing the organization now requires looking beyond the corporate laptop. The modern attack surface has expanded to include unmanaged contractor devices, public-facing cloud assets and the web browser itself, which has become the primary workspace for the enterprise.

As defenders, we face a dual challenge. We must rigorously manage the external exposures that attackers constantly scan for, while simultaneously securing the human interface where users interact with data, AI and the open web. To protect this sprawling environment, security must extend its reach from the external edge down to the browser session.

Reduce the attack surface with active exposure management. Unit 42 found that software vulnerabilities accounted for 22% of initial access for incidents this year, underscoring the urgent need to move beyond simple discovery to active risk prioritization. Effective exposure management bridges this gap by creating a complete, continuous inventory of the digital footprint, including the shadow infrastructure and unauthorized AI tools that traditional scans miss.

Crucially, this strategy must filter out the noise, using threat intelligence to prioritize only those assets that are actively being targeted in the wild (such as CISA KEVs) and lack compensating controls. By focusing limited resources on exploitable, business-critical risks, teams can close the window of opportunity before an attacker finds an open door.

Protect the human interface. The browser is the new endpoint and the new corporate desktop. This is where employees access data, where contractors perform their work and unfortunately, where social engineering attacks like phishing are most effective.

Securing this interface requires an enterprise-grade secure browser that establishes a fully isolated and secured corporate workspace for both managed and unmanaged devices. This powerful layer enforces data controls in real-time, regardless of the underlying hardware. It can disable copy and paste on sensitive pages, prevent file downloads from unknown sources and identify advanced phishing sites that evade standard email filters. By hardening the browser, organizations gain granular visibility into shadow AI usage and directly prevent sensitive corporate data from leaking into unauthorized GenAI tools.

Secure third-party and unmanaged access. The rigid model of shipping corporate laptops to every contractor or acquisition target is no longer sustainable or secure. Organizations need a way to enforce zero trust access on unmanaged devices without the cost and complexity of legacy virtual desktop infrastructure (VDI) solutions.

By securing the workspace through the browser, companies can grant contractors and BYOD users secure access to corporate applications while keeping business data strictly isolated from personal environments. This approach accelerates merger and acquisition integration, and contractor onboarding while ensuring that a compromised personal device cannot be used as a stepping stone into the corporate network.

Collect unified telemetry and automate response. For the endpoints you do manage, data is the fuel for defense. Detecting sophisticated attacks depends on collecting high-fidelity telemetry across processes, network connections and identity behavior, then unifying that data within a central platform.

When this data is analyzed by AI-driven engines, anomalies that would be invisible in isolation become clear indicators of compromise. However, detection is only half the battle.

To minimize damage, response mechanisms must be automated. Security teams must be empowered to isolate compromised endpoints, initiate forensic scans and remediate threats at machine speed, ensuring that a localized infection does not become a systemic breach.

By securing the browser as the primary workspace and rigorously managing the external attack surface, you protect the users and assets that traditional endpoint controls can no longer reach.

05

SECTION 5

Appendix

We organized the data in this section in three dimensions, providing defenders with a clearer view of the patterns we have observed in 2025. First, we outline the MITRE ATT&CK® techniques most closely linked to each tactic. We then present regional and industry-level views that show how investigation types shift across geographies and sectors.

5.1.

OVERVIEW OF OBSERVED MITRE TECHNIQUES BY TACTIC

The following series of charts (Figures 7-18) show the MITRE ATT&CK® techniques we observed in association with specific tactics. Note that the percentages shown represent the prevalence of each technique when compared across the other kinds of techniques identified for each respective tactic. These percentages don't represent how often the techniques showed up in cases (see the website version to explore data about unique techniques and cases).




Figure 7. Relative prevalence of techniques observed in association with the initial access tactic.




Figure 9. Relative prevalence of techniques observed in association with the execution tactic.




Figure 8. Relative prevalence of techniques observed in association with the discovery tactic.




Figure 10. Relative prevalence of techniques observed in association with the persistence tactic.




Figure 11. Relative prevalence of techniques observed in association with the privilege escalation tactic.




Figure 13. Relative prevalence of techniques observed in association with the credential access tactic.




Figure 12. Relative prevalence of techniques observed in association with the defense evasion tactic.




Figure 14. Relative prevalence of techniques observed in association with the lateral movement tactic.




Figure 15. Relative prevalence of techniques observed in association with the collection tactic.




Figure 17. Relative prevalence of techniques observed in association with the exfiltration tactic.




Figure 16. Relative prevalence of techniques observed in association with the command and control tactic.




Figure 18. Relative prevalence of techniques observed in association with the impact tactic.

5.2.

INVESTIGATION TYPE BY REGION

Figures 19 - 21 provide a regional and industry-level view of the investigations handled by Unit 42 during 2025. They show how incident types vary across North America, Europe, the Middle East, Africa and Asia Pacific, alongside a breakdown of the most common investigation categories within the industries most represented in our data. These insights will help leaders understand where activity is concentrated and how exposure differs across sectors and geographies.

The geographic data highlights differences in investigation types regionally, while the industry charts show clear patterns in how threat activity aligns with sector-specific operations and technology stacks. High technology, manufacturing, financial services and healthcare each exhibit distinct mixes of intrusion types, reflecting variation in attack surface, identity architecture and cloud maturity. Together, these views give security leaders a clearer picture of where threats are most active and how the operational context shapes the intrusions Unit 42 investigates.




Figure 19. Investigation type by region: North America.




Figure 20. Investigation type by region: Europe, the Middle East and Africa.




Figure 21. Investigation type by region: Asia-Pacific region.

5.3.

INVESTIGATION TYPE BY INDUSTRY

Figures 22–28 below show a breakdown of the top investigation types associated with the industries most represented in our incident response data, ordered from the industries with the highest volume of investigations to those with lower representation.




Figure 22. Investigation type by industry: High Technology.




Figure 23. Investigation type by industry: Manufacturing.




Figure 24. Investigation type by industry: Professional and Legal Services.




Figure 25. Investigation type by industry: Wholesale and Retail.




Figure 26. Investigation type by industry: Financial Services.




Figure 27. Investigation type by industry: State and Local Government.




Figure 28. Investigation type by industry: Healthcare.



SECTION 6:

Methodology

We sourced data for this report from more than 750 cases Unit 42 responded to between Oct. 1, 2024, and Sept. 30, 2025, with comparisons to metrics from earlier case data going back to 2021. Our clients ranged from small organizations with fewer than 50 personnel to Fortune 500, Global 2000 and government organizations with more than 100,000 employees.

The affected organizations were headquartered in over 50 countries. About 65% of the targeted organizations in these cases were headquartered in the U.S. Cases involving organizations based in Europe, the Middle East and Asia Pacific formed the remaining 35% of the work. Attacks frequently had an impact beyond the locations where the organizations were headquartered.

We combined this case data with insights from our threat research, which is informed by product telemetry, observations from dark web leak sites and other open-source information. Incident responders also contributed their firsthand observations of key trends based on direct work with clients.

Several factors may influence the nature of our dataset, including a trend toward engagements with larger organizations with more mature security postures. We also emphasized cases that reveal emerging trends, which for some topics means focusing on smaller segments of the overall dataset.

For some analysis areas, we chose to filter our data to avoid skewed results. For example, we offered incident response support to help customers investigate potential impacts of CVEs 2024-0012 and 2024-3400, which caused these vulnerabilities to be overrepresented in our dataset. Where appropriate, we corrected for this overrepresentation.

Our guiding principle throughout has been to provide insights into the present and emerging threat landscape, enabling defenders to strengthen their posture against what attackers are doing now and where they are headed next.

Contributors

Amelia Albanese
Sheida Azimi
Jim Barber
Maxfield Barker
Jeremy Brown
Mark Burns
Josh Costa
Kasey Cross
Michael Diakiwski
Dan O'Day
Richard Emerson
Elizabeth Farabee
Robert Falcone
Byrne Ghavalas
Wyatt Gibson

Alexis Godwin
Evan Gordenker
Daniel Gott
Evan Harrington
Tim Heraldo
Brandon Hicks
Manisha Hirani
Jack Hughes
Margaret Kelley
Seth Lacy
Samantha Le
Yang Liang
Chia Hui Mah
Mitch Mayne
Eva Mehlert

Vraj Mehta
Danny Milrad
Jacqui Morgan
David Moulton
Lysa Myers
Erica Naone
Aisling O'Suilleabhain
Aryn Pedowitz
Andy Piazza
Nicholas Pockl-Deen
Brendan Powers
Nathaniel Quist
Adam Robbie
Laury Rodriguez
Sam Rubin

Doel Santos
Mike Savitz
Andrew Scott
Steve Scott
Ram Shenoy
Michael Sikorski
Scott Simkin
Ray Spera
Samantha Stallings
Jenine Sussman
Virginia Tran
Amy Wagman
JL Watkins
Kyle Wilhoit

About Palo Alto Networks

Palo Alto Networks® is the world's cybersecurity leader. We innovate to outpace cyberthreats, so organizations can embrace technology with confidence. We provide next-gen cybersecurity to thousands of customers globally, across all sectors. Our best-in-class cybersecurity platforms and services are backed by industry-leading threat intelligence and strengthened by state-of-the-art automation. Whether deploying our products to enable the Zero Trust Enterprise, responding to a security incident, or partnering to deliver better security outcomes through a world-class partner ecosystem, we're committed to helping ensure each day is safer than the one before. It's what makes us the cybersecurity partner of choice.

At Palo Alto Networks, we're committed to bringing together the very best people in service of our mission, so we're also proud to be the cybersecurity workplace of choice, recognized among Newsweek's Most Loved Workplaces (2023, 2022, 2021), with a score of 100 on the Disability Equality Index (2023, 2022), and HRC Best Places for LGBTQ equality (2022). For more information, visit www.paloaltonetworks.com.

About Unit 42

Palo Alto Networks® Unit 42® brings together world-renowned threat researchers, elite incident responders and expert security consultants to create an intelligence-driven, response-ready organization that's passionate about helping you proactively manage cyber risk. Together, our team serves as your trusted advisor to help assess and test your security controls against the right threats, transform your security strategy with a threat-informed approach, and respond to incidents in record time so that you get back to business faster. Visit paloaltonetworks.com/unit42.

3000 Tannery Way
Santa Clara, CA 95054

Main +1.408.753.4000
Sales +1.866.320.4788
Support +1.866.898.9087

www.paloaltonetworks.com

© 2026 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks, Inc. A list of our trademarks in the United States and other jurisdictions can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.