



STATE OF CLOUD SECURITY REPORT

2025

Table of Contents

Executive Summary 3

Highlights of the 2025 Survey Findings 4

Introduction 6

The Cloud Footprint 7

 Code Ships Fast, and Velocity Runs with Maturity 8

 Half the Compute Already Lives Here 8

 Public Cloud Is the Default, Even for Sensitive Work 9

 Multicloud from Day One 10

 Architecture: Hybrid by Design 10

Application Security 11

 AI Coding Is Taking Over 11

 Shifting Left Has Been a Challenge 11

 Many Security Issues Still Reach Production 12

 Remediating in Production Is Not a Winning Strategy 12

Data Security 13

 Complexity Drives Data Risk 13

 Manual Discovery Marks a Broken Baseline 14

 Posture Leads as the First Line of Defense 15

 How Data Walks Out the Door 15

Incident Response 17

 Escalating Incidents Define the Threat Landscape 17

 Detection and Containment Keep Pace, Resolution Falls Behind 18

AI in the Cloud 20

 AI Security Starts with the Stack 20

Cloud Security Strategy 23

Recommendations for Securing the Cloud 24

 Optimize Pre-Deploy Security Gates 24

 Reduce Incident Response Fragmentation 24

 Fortify Identity and Permissions Management 25

 Leverage AI Security for Proactive Defense 25

 Improve Automation and Remediation Cycles 25

 Extend Cloud Security Operations into the SOC 25

How Palo Alto Networks Helps 26

Methodology 27

Executive Summary

We're witnessing an evolution of cyberattack methods, a wholesale transformation in how adversaries operate. The Palo Alto Networks Unit 42® research team, a leader in threat intelligence, incident response, and proactive services, has documented the shifting threat landscape. The daily cyberattacks we've seen have surged from as much as 2.3 million to up to nearly 9 million in the span of a year, an almost threefold increase driven by attackers' adoption of AI tools. But the volume tells only part of the story.

Concern lies foremost in velocity. Attackers have redefined their key performance indicators, dramatically reducing what we call "mean time to compromise" and "mean time to exfiltrate." Unit 42 testing has demonstrated that breaches that took an average of 44 days in 2021, with AI assistance, can now occur in as little as 25 minutes. The speed, scale, and sophistication we've observed over the past couple of years is incredible.

At the center of this accelerated threat activity sits cloud environments. The fifth annual *State of Cloud Security Report 2025* confirms that more than half of production workloads run in cloud infrastructures, most of which reside in the public cloud. Organizations use an average of six cloud providers, and most DevOps teams deploy new or updated code weekly. Each of these realities expands the potential blast radius of cyberattacks, while compressing the time security teams have to detect and respond.

And the gap between detection and remediation continues to widen. We see this among the one in five organizations reporting that over a quarter of their high or critical issues remain in production

longer than 30 days. While attackers measure breach success in minutes, defenders are still measuring cleanup in weeks.

The rise of large language models (LLMs) and agentic AI pushes the attack surface beyond traditional infrastructure. Adversaries target the tools and LLM systems, the underlying infrastructure supporting model development, the actions these systems take, and critically, their memory stores. Each represents a potential point of compromise.

Our defensive posture must outpace this reality. The foundation remains paramount—hardened identities, segmented access, locked down secrets, and enforced known-good behavior, all with governed dependency hygiene. Now, cut the time it takes to recognize, prioritize, and act by building on that core with AI-enhanced detection and response. Reserve room to probe emerging blind spots, especially around agentic behavior and memory abuse, where yesterday's controls don't reach.

The key mindset is one of layered progression. Attacks are only growing faster and more sophisticated. Staying ahead means focusing on real-time and autonomous security, leveraging AI and automation to secure at the speed of the machine while honing your zero trust architecture and existing controls.

The outcome? A future-proofed AI journey with a secure cloud-native foundation.




Haider Pasha

Vice President, Chief Security Officer, EMEA
Palo Alto Networks


Highlights of the 2025 Survey Findings

Code changes remain in motion.




50+%

of survey respondents deploy new or updated code weekly.



17%

ship daily or faster.



99%

now use GenAI for coding support, and the cycle shows no signs of slowing.

Unit 42 Closeup: Code Changes Remain in Motion


Using LLM coding assistants, while boosting productivity, introduces security flaws into the high-velocity cloud development pipeline. The principal threat lies in generating insecure code and configurations, which attack vectors can facilitate by using context attachment misuse, harmful content generation, direct model invocation, or indirect prompt injection (IPI).

IPI allows malicious instructions, hidden within the external data the LLM processes, to compromise the output, leading to misconfigurations or vulnerable API interactions in the generated code. These LLM-introduced flaws challenge the ability of security teams to enforce standards before deployment.

[Read more](#)




AI is no longer a theoretical risk.



75%

of organizations run in AI production.



99%

have experienced at least one attack on an AI system in the past year.

Unit 42 Closeup: AI Is No Longer a Theoretical Risk

Unit 42 research substantiates AI risk by identifying novel techniques, like Agent Session Smuggling in agent-to-agent systems and advanced Prompt Attacks as key vectors. Furthermore, AI is accelerating the scale and realism of social engineering—the top initial access vector in 36% of incidents that Unit 42 responded to from January 2025 to June 2025. This finding confirms that the AI agent and human user are both critical high-risk attack surfaces.

[Read more](#)



APIs now top the target list.

41%



increase in the amount of attacks on APIs year over year, which is the steepest surge of any threat vector, driven in part by the rapid adoption of AI agents and ungoverned interface sprawl.

Unit 42 Closeup: APIs Now Top the Target List

Flaws in the development pipeline exacerbate the API threat vector. Unit 42 researchers identified critical security risks tied to OpenID Connect (OIDC) misconfigurations within continuous integration/continuous deployment (CI/CD) environments.

Three specific, advanced threat vectors target these weaknesses:

- **Loosely configured policies:** Federation rules fail to enforce meaningful validation on OIDC token claims, granting unauthorized access.

- **Reliance on user-controllable claims:** Attackers manipulate claims in the OIDC token that are designed to be set or influenced by the user to inject malicious values.
- **Poisoned pipeline execution (PPE) synergy:** PPE vulnerabilities combined with permissive identity federation settings allow attackers to leverage highly privileged CI systems to gain broad access to downstream cloud resources.


Read more



The cloud-SOC divide is closing fast.

Unit 42 Closeup: The Cloud-SOC Divide Is Closing Fast

Unit 42 Incident Response supports this strategic imperative, finding that 70% of security incidents now span three or more attack surfaces. This multivector reality demands a single, unified security response model across endpoints, networks, and cloud environments.



of organizations say cloud and application security should integrate with the SOC in a shift that marks the end of siloed control and the rise of unified operations.

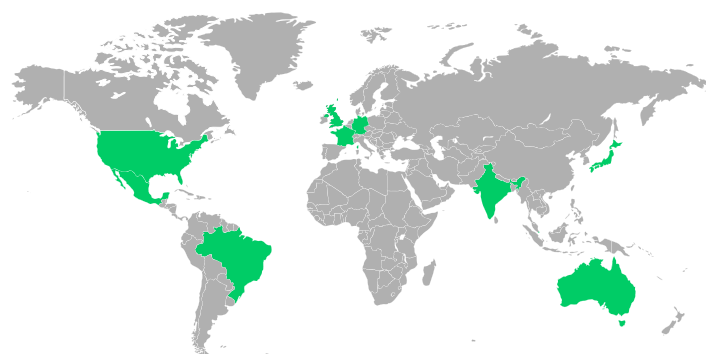
Introduction

Our research included major industry sectors, with representation from financial services, healthcare, technology, media, telecommunications, energy resources and industrials, and consumer products and services. Among the organizations surveyed, 50% came from enterprises with over \$1 billion in annual revenue, and 50% came from organizations with an annual revenue of \$100 million to \$1 billion.

Survey participants included an equal mix of executive leadership and practitioner-level roles to cover a broad spectrum of viewpoints across organizations. Practitioner-level participants were from functions within development, AppSec, cloud security, and security operations. All respondents were sourced from professional survey panels and self-reported as knowledgeable and familiar with their organization's cloud operations and cloud security.

Palo Alto Networks partnered with Wakefield Research, who conducted our survey from September 29, 2025, to October 17, 2025. Wakefield Research gathered data from more than 2,800 respondents in 10 countries: Australia, Brazil, France, Germany, India, Japan, Mexico, Singapore, the United Kingdom, and the United States.

The *State of Cloud Security Report* examines the security practices, tools, and technologies that organizations worldwide are employing to take advantage of cloud services and new application tech stacks.



Wakefield Research gathered data from more than **2,800 respondents in **10 countries**.**

The Cloud Footprint

In the five-year history of this report, fully cloud-native organizations were unheard of—until now. This year's survey found that 7 out of 10 organizations today operate in extensive cloud integration or fully cloud-native states. While the remaining 39% still run on basic infrastructure or limited projects, the balance has clearly tipped.

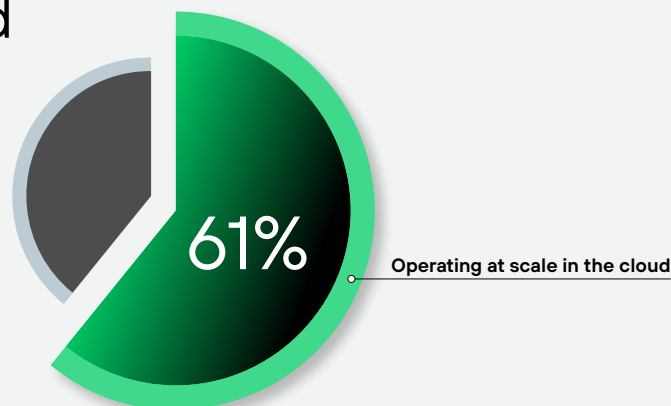

Two plateaus define the current landscape. Among survey participants, 38% describe themselves as extensively integrated, meaning adoption of multiple cloud service providers (CSPs), optimized resource management, and early container or cloud-native use. Another 23% report fully cloud-native operations with high automation and continuous delivery, which represents nearly one in four organizations already functioning at enterprise scale.

Organizations manage environments that span an average of six cloud providers, alongside a layered mix of IaaS, PaaS, and SaaS platforms—a configuration that reflects how hybrid architectures have become the default state of enterprise IT.

Enterprise cloud has entered a scaled, high-velocity, multiprovider phase, where public cloud leads, sensitive data is distributed across environments, and mixed runtimes are the norm.

For security leaders, the implication is straightforward. Cloud risk is defined by the operational complexity of mature, multiprovider ecosystems, not by migration readiness or basic posture management. The challenge has shifted from organizations getting to the cloud, and once they get there, to maintaining consistent visibility, policy, and control.

How Organizations Adopt Cloud




Code Ships Fast, and Velocity Runs with Maturity

Among the organizations surveyed, 53% deploy new or updated code to production at least weekly. The cadence rises to 57% among the mature cloud cohort and sits statistically higher than 45%, which early adopters reported. Daily or more frequent releases reach 17% overall and 19% inside the extensive or fully cloud-native group, confirming high-frequency deployment inside large, multiprovider enterprises.

Release velocity defines the response window for security and compliance. Weekly and daily pipelines demand controls that operate at runtime speed, such as policy gates in CI, automated checks across infrastructure as code (IaC) and dependencies, identity and entitlement verification, and continuous posture evaluation that feeds back into the pipeline. Quarterly review cycles can't govern systems that change hundreds of times between meetings.

Speed to Production




Half the Compute Already Lives Here


Survey results indicate that an average of 51% of organizations run workloads in the cloud. Organizations that describe themselves as extensive or fully cloud-native run higher at 52%, compared to early cloud adopters at 48%.

The distribution shows depth, where 48% place 50%–74% of their workloads in the cloud, and another 12% report placing 75%–100% of their workloads in the cloud. A small segment of 11% falls below 25%. Enterprise estates now lean majority cloud.

Regional standouts include Mexico at 54% and Japan at 55%, while Singapore and the UK report 47% and 48%, respectively.




Cloud Saturation Worldwide



Public Cloud Is the Default, Even for Sensitive Work

Respondents report that 55% of workloads run in a public cloud IaaS or PaaS. The average rises to 57% in the fully cloud-native cohort and is statistically higher than early-stage organizations. Private cloud and on-premises remain significant in the mid-40% range, yet public cloud now edges them out.

Where Sensitive Data Lives



Sensitive data follows a similar pattern, distributed across multiple locations. Roughly one quarter remains in owned on-premises infrastructures, with statistically similar mid-20% shares in public cloud and SaaS storage, and endpoints carrying a double-digit share.

Multicloud from Day One

Among respondents, 65% report using between three and nine CSPs. Another 7% report they use 10 or more CSPs. Multicloud spans maturity levels and company sizes. Organizations with limited cloud projects still report an average of six providers, which is statistically similar to averages in more mature cohorts.


Multicloud now represents the starting condition for security architecture. Most teams inherit a provider mix that often spans half a dozen vendors before they claim full cloud-native status.

Sensitive data, regulated workloads, and AI-assisted development all now live in the same runtime.

Architecture: Hybrid by Design

No single runtime model defines the cloud. Respondents report that 26% of their cloud workloads run on virtual machines, which is still the largest individual share. Both self-hosted and managed containers account for a combined 40%, while PaaS and serverless architectures make up the remaining third. The compute layer today isn't transitional. Teams are running lift-and-shift workloads, containerized services, and fully abstracted functions side by side.

The Architecture Breakdown




Application Security

AI Coding Is Taking Over

Generative AI is raising the stakes for security teams. Among organizations surveyed, 99% are using GenAI tools to assist in software development. They introduce high-volume streams of AI-generated code into pipelines that the speed and complexity of cloud-native development already strain.


Shifting Left Has Been a Challenge

While logically it has made sense to “shift security left”—allowing developers to secure their code as they go—the idea of putting in gating mechanisms to prevent insecure code from reaching production hasn’t been well received. The survey found that 85% of respondents find security to be a hindrance to delivering software releases.

99% 
of organizations
use GenAI tools to
support coding.

AI is turning AppSec on its head. Supercharged development pipelines are leaving growing backlogs.

How Often Security Delays Major Releases



Many Security Issues Still Reach Production

While the cloud has enabled greater development speed, security issues now regularly reach production.

The survey indicates that 53% of organizations now deploy new code weekly but struggle to stop security issues from reaching production. Only 34% of organizations can prevent all but 10% of high and critical security issues. One in five teams let upwards of 37% of high and critical issues into production.

When asked why more issues aren't prevented from reaching production, the top answers included concerns about slowing development velocity and the limited ability to integrate tools into CI/CD pipelines.

The remaining responses—false positives (18%) and developer resistance (18%)—compound the problem, showing respondents feel that, if they attempt stronger controls, they face accuracy or adoption challenges. Moreover, 53% of organizations report blocking based on severity alone or don't block all.

Remediating in Production Is Not a Winning Strategy

Fixing code issues remains a slow process for most teams. Across organizations in the survey, 82% report that it takes longer than a week to deploy a code fix in production.

Prioritizing can be another challenge for organizations that can't contextualize their risks. Only 9% of organizations are using runtime context as their primary way to prioritize risks—leaving it unclear whether risks are exploitable or running in production.

Organizations are under water with remediating issues. Meanwhile, applications persistently carry risks and backlogs grow larger.

53%

of organizations
now deploy new
or updated code to
production weekly.

82%

of organizations
report that it takes
longer than a week to
deploy a code fix
in production.

Why Organizations Avoid Enabling Stronger Pre-Merge or Pre-Deploy Guardrails




Data Security

Complexity Drives Data Risk

Teams agree on the source of the problem. Data risk grows out of complexity. The survey found that 60% of organizations cite fragmented cloud environments as their top data security challenge, a finding consistent across executive and practitioner levels, company sizes, and industries. The only regional variance appears in the UK (63%), Australia (64%), and Singapore (64%), where organizations feel the challenge more acutely, perhaps due to regulations like GDPR and PDPA. Cloud growth carries a cost, where, as adoption accelerates, architectural sprawl follows and visibility fractures.

Cloud data security has shifted from a compliance exercise to an architectural challenge, where fragmented environments, expanding identities, and uncontrolled data flows define the primary fault lines of enterprise risk.

The Data Security Struggle



IAM Remains the Weakest Link

Identity is the next fault line after complexity. Across survey participants, 53% point to lenient identity and access control management (IAM) practices or insufficient permission granularity as a top challenge. The number climbs to 57% among organizations operating more than six AppSec tools. Identity exposure scales with tooling complexity, with each additional control plane widening the access surface and eroding the ability to maintain least privilege.

Secrets Still Slip Through

Issues with handling secrets are far from solved. According to 45% of respondents, poor secret-management practices are an ongoing struggle, a number that rises to 47% among SOC teams and organizations with less time in the cloud. These findings reinforce a familiar refrain. As environments scale and diversify, traditional boundary-based controls can't keep up. Secrets proliferate across pipelines, workloads, and integrations faster than static scanners or manual audits can track.

Manual Discovery Marks a Broken Baseline

Among all organizations in the survey, 48% still rely on manual review to identify and classify sensitive data. At cloud scale, this approach breaks down.

As noted previously, with 60% of respondents citing cloud complexity as the top data security challenge, fragmentation explains much of the problem. Each new SaaS platform, unmanaged data flow, and ephemeral cloud asset increases the likelihood that sensitive data ends up in an untracked location. Security teams fall back on manual discovery because it's the only option when visibility fails. They're not reviewing files; rather they're hunting for repositories they didn't know existed.

The burden of manual review reflects years of technical debt. Without consistent tagging, enforced standards, or automated inventory, teams can't trust their coverage. Until discovery matures and scales, manual review remains the default response to a sprawling, unmappable data landscape.

48%

of organizations still
rely on manual review
of sensitive data.


60%

of organizations
cite fragmented
cloud environments
as their top data
security challenge.

Posture Leads as the First Line of Defense

The most consequential data exposures in the past year were first detected by data security posture management (DSPM) at 31%, followed by EDR/XDR at 26%, and firewall at 25%. DSPM's lead in detecting cloud data misconfigurations reflects a growing reliance on cloud-native controls to surface risk earlier.

Which Controls Spot the Breach



How Data Walks Out the Door

As the survey found, SaaS sync or export misuse lead as data exfiltration vectors at 63%, followed by overpermissive external sharing (59%) and compromised credentials or tokens (58%). Misconfigured public access ranks lower at 30%, although it still represents nearly one in three organizations. Large enterprises show a statistically higher rate for misconfigured public access, which is consistent with the configuration drift and baseline-control erosion that accompany scale. And, 28% of respondents report insider transfer to unmanaged endpoints.

Unit 42 Perspective: The Challenge with Tokens

Tokens are often regarded as the invisible currency of trust in cloud environments, enabling seamless automation and interapplication communication. However, poor token management transforms them into a dangerous tool for threat actors. Three recurring security patterns demonstrate this danger:

1. **Dormant integrations**, where forgotten, unused tokens remain active.
2. **Insecure token storage**, where tokens are left unprotected or unencrypted.
3. **No expiration or rotation**, where compromised tokens remain valid indefinitely for prolonged breaches.

Read more [→](#)

The following insights stand out.

1

Data leaves through both legitimate business systems and breach events. The top vector—SaaS sync or export misuse—points to collaboration tools, cloud storage, business intelligence platforms, and AI assistants as primary data-loss channels. The exposure is operational, not exploit-based.

2

Identity and entitlement weaknesses cut across every major vector. Overpermissive external sharing (59%) and compromised credentials or tokens (58%) track almost identically to SaaS misuse, making data exfiltration in the cloud fundamentally an identity problem, not merely a workload or posture problem.

3

Maturity changes the shape of exposure. Organizations that have spent more than five years working in the cloud report higher rates of SaaS misuse (66%) and misconfigured public access (32%), both statistically higher than less-mature peers. Seasoned cloud environments rarely suffer from obvious Amazon S3 open-to-the-world mistakes. Their risks arise from subtler issues, such as persistent oversharing between tenants, token abuse in automation, and uncontrolled synchronization between SaaS systems.



Incident Response

Escalating Incidents Define the Threat Landscape

Every organization surveyed reports experiencing all 10 measured security incidents in the past year. Their response confirms that exposure is more a matter of operating in today's environment than it is missteps.

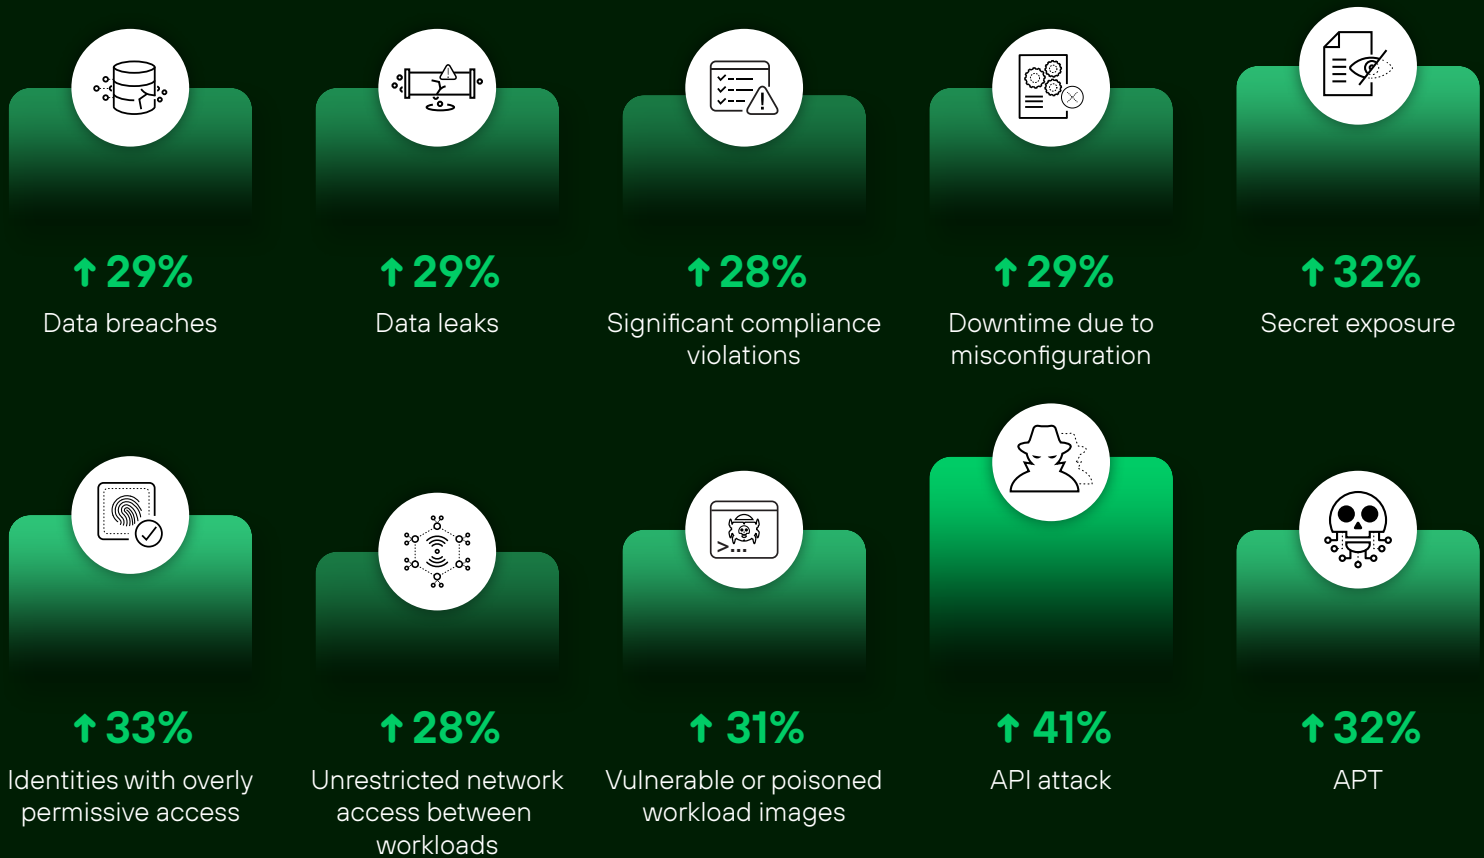
We see the most consequential trend at the intersection of GenAI and API risk. API attacks show the sharpest year-over-year increase at 41%, fueled in part by two compounding forces. First, generative AI has lowered the barrier to exploitation by enabling low-skilled actors to generate high-fidelity attacks. Second, the proliferation of AI agents—many rapidly deployed and lightly governed—has introduced prompt injection vectors and an explosion of API surfaces. Compromise, in other words, now scales with automation and expands through endpoints few teams know they own.

In parallel, advanced persistent threats (APTs) continue to rise, with 32% of organizations seeing an increase in long-term, stealth attacks. The campaigns exploit identity and access misconfigurations, which are also on the rise. Across the organizations surveyed, 33% report increased threats that target overly permissive identities, highlighting the appeal of the control plane as a path to persistent access.

Incident response in the cloud has reached a breaking point, where growing signal volume and tool complexity outpace teams' ability to assemble a single, coherent view of an attack.

Together, these trends show a split-threat model. One branch moves fast, extracting data through high-volume, low-friction API compromise. The other branch embeds slowly, capitalizing on weak governance. Both forms thrive under conditions of complexity, fast-moving parts, and fragmented oversight.


Where Attacks Are Climbing



Detection and Containment Keep Pace, Resolution Falls Behind

When a threat is in progress, 50% of our survey group reports that 50% of security analysts' time goes to data collection and correlation. Another one in five security analysts spend as much as 80% of their time on data correlation. Attackers, meanwhile, exfiltrate data at record speeds in 2025.

Where Analysts Spend Time



What's holding up defenders? The issues, while many, point to communication shortfalls. Among respondents, 50% cite disjointed workflows between cloud and SOC teams, and another 50% point to isolated data sources. The difficulty compounds downstream, where 49% struggle to unify alerts into a coherent incident story, and 42% lack a unified timeline that bridges cloud and enterprise telemetry.


The extent of fragmentation disrupts momentum. Most organizations (74%) detect and contain threats within 24 hours, but too few can carry that velocity through to resolution. One in three teams need more than a day to close an incident. Of these, 9% need between a week and a month. Context loss, unclear ownership, and disconnected systems prevent teams from moving efficiently from insight to action.

Incident response today suffers from an absence of cohesion. While teams see plenty, they can't see it together. Until incident response architectures collapse the divide between posture, runtime, and response, fast detection won't close the loop.

From Detect to All Clear



33% of teams need **a day to a week** to resolve an incident.



Of these, 9% need **a week to a month or more**.

33% 
of organizations
need a full day or
more to resolve a
security incident.

Unit 42 Perspective: Incident Response Process

In 2024, 29% of incidents that Unit 42 responded to involved cloud or SaaS environments. Cloud investigations require a shift in mindset compared to traditional, endpoint-focused forensics. Unit 42 outlines a five-step incident response process tailored for cloud environments, focusing on investigating identities, misconfigurations, and service interactions, as well as key steps to strengthen cloud defenses and increase resiliency.

Read more 

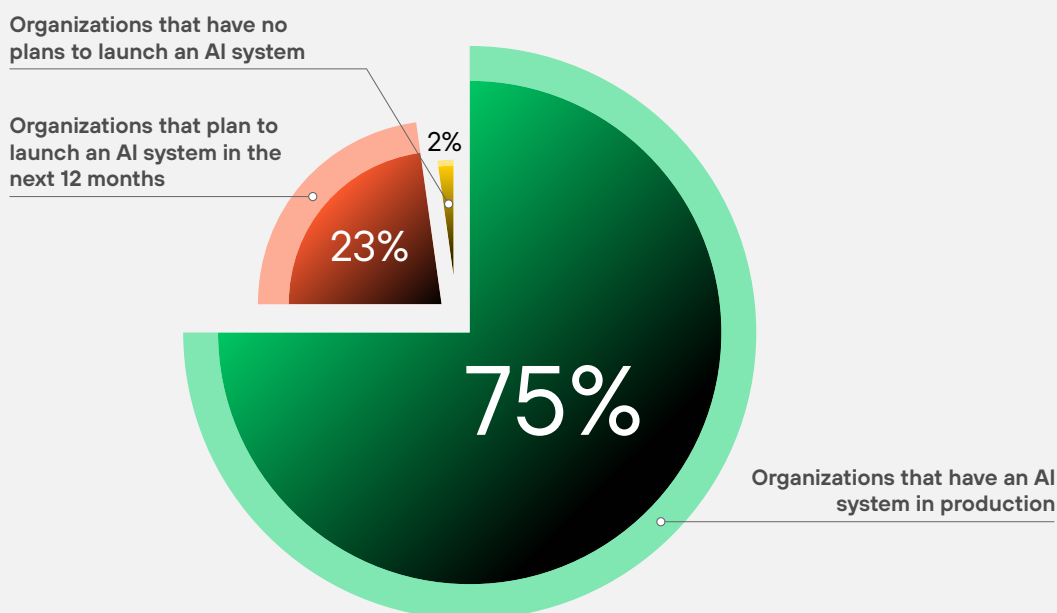
AI in the Cloud

AI Security Starts with the Stack

AI systems are already in production at scale. Of the organizations in the survey, 75% have deployed AI to the cloud, with another 23% planning to within the next 12 months. What's notable is where defenders feel the most exposed and what they've already seen.

AI is no longer confined to R&D or pilot programs. Most organizations now run production AI systems, embedding new models and workflows directly into the cloud infrastructure they already depend on.


Organizations' Relationship with AI



Top AI Security Concerns, Ranked

1

Securing the cloud infrastructure and CI/CD pipelines



2

Protecting sensitive training data



3

Ensuring compliance with new AI regulations




4

Managing risks from open-source AI libraries



5

Securing the AI models themselves



Asked to name the top security concern across the AI/ML development lifecycle, leaders don't point to the model or prompt, not initially. They first point to the environment. The top concern for 26% of leaders is the underlying cloud infrastructure and CI/CD pipelines, followed by protection of sensitive training data (20%), and compliance with emerging regulation (19%). The attack surface, it turns out, hasn't moved far. It's still grounded in cloud infrastructure.



Attack activity justifies this posture, with 99% of organizations having experienced at least one attack on an AI system in the past year. The most common breach path—reported by 47%—involves data exfiltration through assistants or plugins. Almost as many organizations point to model supply chain tampering, model endpoint abuse, and token theft. Prompt injection and output manipulation follow at 43%. Nearly all of these threats involve an API boundary, reinforcing the role of ungoverned interfaces in a scalable AI compromise.

Taken together, the threat landscape is less about model compromise than about risks scaled at speed—exposed endpoints, manipulated inputs, poisoned packages, overpermissive access, and runtime blind spots—now running in AI-accelerated pipelines with little margin for delay or error.


Unit 42 Perspective: Poisoned Packages

One example of the dangers of poisoned or malicious packages is the recent npm supply chain attack. The attack began with a credential-harvesting phishing campaign that spoofed npm's MFA login to gain initial access to developer accounts. The threat actor then deployed a malicious worm within compromised package versions. Unit 42 assesses that an LLM based on code features might have assisted the malicious code. When installed, the worm's script would execute, scanning the developer's environment for sensitive credentials such as SSH, cloud, and cryptocurrency keys. It created a multistage supply chain compromise that enabled the malware to propagate itself under different package names.

[Read more](#)



AI Attacks in the Wild



Cloud Security Strategy


Security leaders are done chasing marginal gains from tool accumulation. On average, organizations report using 17 security tools from five different vendors to manage risk. Nearly all organizations (97%) prioritize reducing that footprint, with nine in 10 favoring a centralized security solution that sits across all cloud accounts and services. What's more, 89% want a platform solution that integrates cloud and application security with the SOC.

The most significant finding, though, isn't about platform reach. In what can only be considered a strategic realignment, 89% of respondents say cloud security and security operations should merge.

As cloud adoption matures, security teams reevaluate the structure, focus, and effectiveness of their entire toolchain.


While early cloud adoption drove a separation between engineering-led cloud security and SOC-driven detection and response, the coming of age or maturation of cloud adoption, along with the nature of modern attacks, are pressing for change. Security practitioners and executives alike see cloud as central to the operational model.

What Security Teams Want Now



of organizations prioritize reducing tool accumulation.

9 in 10



of respondents say cloud security and security operations should merge.

Recommendations for Securing the Cloud


Optimize Pre-Deploy Security Gates

Enforce targeted, context-aware security checks that prioritize exploitability and business impact, as well as Common Vulnerability Scoring System severity. Too many critical issues still reach production because premerge guardrails create noise without signal. Integrate controls directly into CI/CD pipelines and fine-tune them with exploitability context to reduce false positives.

Reduce Incident Response Fragmentation

Cloud risk decisions are slow because organizations can't assemble one timeline, one owner, and one understanding of "urgent." Collapse tool and team silos through a unified investigation platform. Prioritize automated correlation and deduplication using SOAR or equivalent orchestration layers. The platform should deliver a single, chronological timeline across cloud and enterprise telemetry to enable seamless collaboration between AppSec, SOC, and cloud teams. When detection and response live in different systems, MTTR drags.

Security doesn't fail from a lack of intention. It breaks where execution stalls, where workflows fracture, and where responsibility diffuses across teams and tools.



Fortify Identity and Permissions Management

Treat IAM and secret management as a tier-one security priority, enforcing granular, least-privileged access across all cloud environments and third-party applications. Implement tighter permission boundaries and connect permissions to real-time usage context. Eliminate high-risk, inefficient manual reviews to ensure sensitive data is consistently and correctly identified and protected at cloud scale. If you haven't integrated DSPM into your security tech stack, consider shifting your budget toward a solution that provides continuous visibility and control over cloud configurations and data access, as it's the most effective detection vector for high-consequence exposures.

Leverage AI Security for Proactive Defense

With 75% of organizations already running AI systems in production, the AI supply chain—model, pipeline, and environment—has become an attack surface. Secure your cloud infrastructure, CI/CD workflows, and data pipelines that support model training and deployment. Hardening the development surface early helps prevent prompt injection, model tampering, and API abuse before they reach production.

Improve Automation and Remediation Cycles

Many high-severity issues—particularly around secrets, misconfigurations, and open-source components—continue to age beyond 30 days. Expand automation across detection, prioritization, and remediation workflows. Focus on high-volume, recurring issues that strain analyst time and increase exposure windows. Automation saves time, enables decisions, and compresses the distance between identification and closure.

Extend Cloud Security Operations into the SOC

Cloud security and enterprise security operations are no longer separable. Unifying cloud telemetry with alert workflows and response actions is a top leadership priority. Collapse posture, detection, and response into a shared system where alerts correlate automatically, incident timelines span all environments, and response actions trigger from a single interface. The security organization needs only one system of action, not more tools.

How Palo Alto Networks Helps

Cortex® Cloud™ is the industry's first agentic-native cloud security solution engineered from day one to unify posture, prevention, and response across the full application lifecycle—from code to cloud to SOC. Built on a single data lake, it normalizes telemetry for full-context, AI-driven prioritization, and automated remediation to stop threats in real time.

Unlike legacy cloud-native application protection platforms (CNAPPs) that operate in data silos, Cortex Cloud fuses CSPM, DSPM, AI-SPM, CWPP, ASPM, CIEM, and CDR into a platform that integrates with DevOps and SOC workflows alike. Teams gain continuous visibility into their cloud estate, real-time threat detection aligned to the MITRE ATT&CK® framework, and the ability to respond from the same environment where risk is identified.

To further enhance how teams respond to risk, Cortex Cloud is now integrated with the industry's most capable agentic AI built into a cloud security platform. Advanced AI threats require a brand-new approach. It requires an agentic AI workforce trained on 1.2 billion real-world responses and secured by human-controlled guardrails to give teams a powerful way to quickly and comprehensively resolve the most complex security challenges.

Security leaders use Cortex Cloud to replace disconnected tools with a coordinated signal, speed up decision-making, and reduce dwell time with built-in automation. Because Cortex Cloud is unified by design, it eliminates redundant data and maintains context across tools to empower teams so they can see what matters and act fast.

The Palo Alto Networks Cortex Cloud platform secures applications from code to cloud to SOC, across multicloud environments.



Methodology

The fifth annual *State of Cloud Security* survey was conducted between September 29, 2025, and October 17, 2025, using an email invitation and an online survey.

The respondent population comprised 2,800 executives and practitioners from development, information security, or information technology departments across four key regions: NAM (US, 36%), EMEA (DE, FR, and UK, 21%), LATAM (BR and MX, 14%), and APJ (IN, SG, AU, and JP, 29%).

This survey also sampled input from all major industries, garnering representation from consumer products and services, energy resources and industrials, technology, media and telecommunication, financial services, and healthcare.

Results of any sample are subject to sampling variation. The magnitude of the variation is measurable and affected by the number of interviews and the level of the percentages expressing the results. For the interviews conducted in this study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 1.9 percentage points for the global sample, 3.1 percentage points in NAM, 4.0 percentage points in EMEA, 4.9 percentage points in LATAM, 3.5 percentage points in JPAC, 6.9 percentage points in each market (Germany, France, UK, Brazil, Mexico, India, Singapore, Australia, and Japan), and 3.1 percentage points in the US from the result that would be obtained if interviews had been conducted with all persons in the universe represented by the sample.

About Palo Alto Networks

As the global cybersecurity leader, Palo Alto Networks (NASDAQ: PANW) is dedicated to protecting our digital way of life via continuous innovation. Trusted by more than 70,000 organizations worldwide, we provide comprehensive AI-powered security solutions across network, cloud, security operations and AI, enhanced by the expertise and threat intelligence of Unit 42. Our focus on platformization allows enterprises to streamline security at scale, ensuring protection fuels innovation. Explore more at www.paloaltonetworks.com.

About Cortex Cloud

Cortex Cloud, the next generation of Prisma® Cloud, merges best-in-class CDR with industry-leading CNAPP for real-time cloud security. Harness the power of AI and automation to prioritize risks with runtime context, enable remediation at scale, and stop attacks as they occur. Bring together your cloud and SOC on the unified Cortex platform to transform end-to-end operations. Experience the future of real-time cloud security at www.paloaltonetworks.com/cortex/cloud.

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

© 2025 Palo Alto Networks, Inc. A list of our trademarks in the United States and other jurisdictions can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies.



CYBERSECURITY
PARTNER OF CHOICE