2026 RESEARCH REPORT

The State of Cloud Security 2026

Signal vs. Noise: How AI Is Reshaping Cloud Security Operations

What 633 cybersecurity professionals across nine countries say about the gap between AI's promise and operational reality — and what it means for security teams in 2026.

Contents

Executive Summary

Cloud security is no longer a standalone concern — it's the backbone of modern defense. And expectations for AI-powered protection are rising fast. Vendors promise autonomous protection, instant detection, and seamless remediation. But what do security teams actually need?

To cut through the hype, Prowler surveyed 633 cybersecurity professionals across nine countries about what they actually need from cloud security in 2026. The findings reveal a workforce under immense strain, a widening gap between AI ambitions and adoption, and a clear set of priorities that should shape every vendor's roadmap. The question isn't whether AI will reshape cloud security. It's whether it will reduce noise or amplify it.

This report isn't just a snapshot of the market. It's a blueprint for what AI-native cloud security must become: contextual, transparent, community-informed, and operationally grounded.

Key Findings at a Glance

71
Weekly Incidents
Average per team, 3,600+ annually

18%
AI at Scale
Have achieved autonomous AI

46%
Trust Gap
Doubt autonomous AI systems

42%
Skills Shortage
Cite talent gaps as top barrier

53%
Cost Concerns
Say AI is too costly or complex

67%
Want ROI Proof
Need cost savings to adopt

> The core insight: Detection is no longer the hard problem. What's broken is everything that happens after the finding shows up: context gathering, re-triage, lost institutional knowledge, and the manual toil of stitching siloed data together.

Chapter 1

The State of Cloud Security Operations

How weekly incidents, skills shortages, and compliance complexity are overwhelming security teams — and why more tools aren't the answer.

The State of Cloud Security Operations

Security teams are drowning. Not because they lack tools, but because the tools they have generate more work than they resolve. The average security team now handles 71 incidents per week — more than 3,600 every year — and more than a quarter of respondents say they spend over half their time on low-value manual tasks like triaging alerts, gathering context, or assembling compliance evidence.

This isn't a technology problem. It's an operational architecture problem. Security engineers have become, in effect, human integration layers — flipping between fifteen tabs, trying to piece together whether a finding actually matters in the context of their environment. That's not security work. That's data assembly. And it's burning people out.

Consider what a typical cloud security engineer's morning actually looks like. Their CSPM flags a publicly exposed S3 bucket. Their identity governance tool surfaces an IAM role with overly broad AssumeRole permissions. Their container security scanner detects a critical CVE in a base image running in EKS. Their SIEM correlates an anomalous API call pattern from an EC2 instance in a production VPC. Each of these findings lives in a different console, with different severity scoring, different context, and no shared understanding of how they relate to each other.

The engineer has to manually correlate these signals: Does that IAM role have a trust policy that could be chained with the exposed bucket? Is the vulnerable container image actually reachable from the internet, or is it behind a private subnet with restrictive security groups and NACLs? Is the anomalous API activity related to the overprivileged role, or is it a legitimate workload pattern? Answering these questions requires toggling between AWS Console, CloudTrail logs, VPC flow logs, Kubernetes RBAC policies, and half a dozen third-party dashboards. The tools find problems in isolation. The human stitches the story together.

The Strain Isn't Distributed Evenly

The survey reveals sharp divides in how teams experience this pressure. Companies with fewer than 100 employees were 33% more likely to automate fewer than 10% of their security incidents. Meanwhile, CISOs were 73% more likely to report automating 50–75% of incidents — showing that automation maturity is being driven from the top, but is failing to trickle down to the teams that need it most.

Smaller organizations remain stuck in manual workflows, creating an efficiency drain for the teams that can least afford it. While larger enterprises have begun investing more heavily in automation, the gap between strategy and execution remains wide across the board.

What's Actually Slowing Teams Down

When asked about their biggest operational challenges, the answers were consistent across geographies and company sizes:

42%	39%	34%
Skills Shortage	**Compliance Burden**	**Limited Automation**
AI and cloud security talent gaps	Growing regulatory complexity	Gaps in detection, triage, remediation

These three challenges form a compounding cycle. Talent shortages mean fewer people to manage growing compliance requirements. Limited automation means those few people spend their time on repetitive manual work. And the resulting burnout accelerates turnover, deepening the skills gap further.

The compliance burden alone has become a major operational tax. Teams are now managing overlapping frameworks such as SOC 2, ISO 27001, PCI DSS, HIPAA, GDPR, and increasingly sector-specific regulations, each requiring evidence collection across cloud resources, access controls, encryption configurations, logging pipelines, and network segmentation. Mapping a single Terraform-managed infrastructure change against multiple compliance frameworks can take hours of manual cross-referencing. When auditors ask for evidence that all production RDS instances enforce encryption at rest with customer-managed KMS keys, someone has to go pull that data, screenshot it, document it, and repeat the process quarterly.

> Every security team has that senior engineer who's been around for years and just knows things: why that weird Lambda permission exists, the remediation playbook for that specific CVE, which team to call when a critical finding needs immediate attention. When that person leaves, all of that knowledge evaporates. It's in Slack threads that are buried. It's in Confluence pages that nobody can find. It's gone.

Chapter 2

The Cloud Platform Landscape

Why familiarity still drives adoption, multi-cloud remains aspirational, and teams want vendors to meet them where they are.

The Cloud Platform Landscape

Cloud security adoption is accelerating, but the platform landscape remains surprisingly concentrated. Despite the industry's multi-cloud ambitions, most teams are gravitating toward the ecosystems they already know and trust.

The Platform Hierarchy

36%
AWS
Primary cloud provider for workloads

28%
Azure
Second-largest cloud share

13%
Google Cloud
Distant third overall

The U.S. market is even more AWS-dominant, with respondents 21% more likely than average to choose AWS as their primary provider. Healthcare stands out as an exception: respondents in that sector were 93% more likely to use Google Cloud Platform, suggesting GCP has carved out meaningful traction in regulated industries.

Multi-Cloud: Aspiration vs. Reality

Despite the industry buzz around multi-cloud strategies, only 16% of organizations are operating in true multi-cloud environments today. For most teams, consolidation — not expansion — remains the norm.

This preference reflects operational reality. Teams want vendors to meet them where they are, not force technology resets that drain resources and slow progress. In 2026, success will favor solutions that work within established ecosystems rather than demanding teams rebuild from scratch.

> The takeaway for security leaders: cloud security expectations are shaped largely by what teams already know and trust. Vendors that integrate seamlessly with existing infrastructure will win. Those that require wholesale platform migration will lose.

Chapter 3

The AI Adoption Gap

Why ambitions outpace adoption, trust remains the biggest barrier, and only 18% have achieved autonomous AI at scale.

The AI Adoption Gap

AI permeates every corner of the cloud security conversation, yet its practical maturity is still catching up to the promise. The survey reveals a market that's eager but cautious, with a significant gap between what teams want AI to do and what they've been able to operationalize.

Where Teams Actually Are

31%
Exploring
Still in early experimentation

18%
Autonomous at Scale
Fully deployed AI systems

51%
Agentic AI Adopted
Integrated into cloud operations

Companies with fewer than 100 employees were 50% more likely to still be in the early exploration phase. CISOs, by contrast, were 46% more likely to report scaled, autonomous AI use — confirming that AI adoption is being pushed from the top down. But implementation remains uneven across organizational layers, with frontline teams often lacking the training, budget, or mandate to fully operationalize AI capabilities.

The Trust Barrier

The biggest blockers to autonomous AI adoption aren't technical — they're organizational:

53%
Cost & Complexity
AI solutions too expensive or complex

46%
Lack of Trust
Doubt in autonomous AI systems

25%
Integration Friction
Poor multi-cloud tool integration

The trust deficit is particularly significant. Security teams are being asked to hand over access to their cloud environments, their findings, and their risk decisions to AI systems. That's a massive trust question. Nearly half of all respondents aren't convinced these systems can be relied upon to act autonomously.

This trust gap creates an opening for vendors who can demonstrate transparency, auditability, and a track record of accuracy. Security teams don't just want AI that works — they want AI they can verify, challenge, and understand.

> There’s a meaningful difference between bolting a chatbot onto a dashboard and building an AI system that’s grounded in years of practitioner knowledge, that understands enterprise context, that learns from your organization’s history, and that was built by people who’ve been living and breathing cloud security since before it was a market category.

PROWLER

Chapter 4

The Visibility Crisis

Why teams are seeing everything but understanding nothing — and how attack path visualization changes the game.

The Visibility Crisis

Ask any security practitioner what they need most, and the answer almost always starts with visibility. But the survey reveals a paradox: teams have more data than ever, yet feel less confident in their ability to act on it.

Seeing Everything, Understanding Nothing

31% of respondents rate their cloud visibility as only “average,” and just 30% feel confident in their ability to respond to threats in real time. Companies with more than 5,000 employees were 43% more likely to say they are not very confident in their real-time detection capabilities — suggesting that scale creates its own visibility challenges.

Limited visibility slows detection, erodes trust in automation, increases analyst burnout, and makes it harder to assess risk with any confidence. When teams can’t distinguish signal from noise, every alert becomes suspect and decision-making grinds to a halt.

The Real Problem: Context, Not Data

The issue isn’t that teams lack telemetry — most environments are generating CloudTrail logs, VPC flow logs, GuardDuty findings, Config rules, and container runtime signals around the clock. The problem is that these signals exist in isolation, with no connective tissue between them.

Attack Path Visualization: Seeing How Threats Chain Together

One of the most critical visibility gaps the survey exposes is the inability to see attack paths — the chains of misconfigurations, excessive permissions, and network exposures that an attacker could traverse to reach high-value assets. Individual findings are easy to generate. Understanding how they combine into exploitable paths is where most tooling falls short.

Consider a real-world scenario: A security scanner identifies an EC2 instance with an overly permissive instance profile that grants s3:* permissions. In isolation, that's a medium-severity finding. But when you layer in context — the instance is in a public subnet with a security group allowing inbound SSH from 0.0.0.0/0, it's running an unpatched Apache version with a known RCE vulnerability, and the S3 buckets it can access contain PII subject to GDPR — you're looking at a complete attack path from the internet to sensitive data with no authentication barriers.

Without attack path visualization, each of these findings gets triaged independently. The SSH exposure might be deprioritized because the team has a bastion host policy. The unpatched Apache instance might sit in a remediation queue. The S3 permissions might be flagged but accepted because the instance "needs" that access. No single finding screams critical. But the combination is a breach waiting to happen.

Effective attack path analysis requires stitching together data from multiple domains: IAM policies and trust relationships, network topology including VPC peering, Transit Gateway routes, and security group chains, resource configurations, vulnerability scan results, and runtime behavior signals. Only 17% of survey respondents prioritized unified visibility as a capability they want from AI — not because it's unimportant, but likely because most teams haven't experienced what true cross-domain visibility looks like. Those that have understand it changes everything about how you prioritize.

The IAM Complexity Problem

Identity and access management has become the most complex attack surface in cloud environments, and it's where the visibility crisis hits hardest. AWS IAM alone involves the interaction of identity policies, resource policies, permission boundaries, service control policies, session policies, and VPC endpoint policies. A single API call might be evaluated against six or more policy layers before it's allowed or denied.

When a scanner flags an IAM role as overly permissive, the finding is technically correct but operationally useless without understanding the full effective permissions picture. That role might have a broad identity policy, but an SCP at the organizational unit level restricts it to specific regions and services. Or the resource it accesses might have a bucket policy that limits access to a specific VPC endpoint. Determining effective permissions requires resolving all of these layers simultaneously — a task that takes experienced engineers 30 minutes or more per role, and most organizations have hundreds or thousands of roles to evaluate.

The Re-Investigation Trap

Consider a common scenario: a scanner flags an overly permissive IAM role on a production workload. Red, critical, top of the queue. But what the tool doesn't know is that three weeks ago, the team reviewed that exact finding, determined there's a compensating control — perhaps a service control policy at the org level, perhaps network-level restrictions — and the CISO explicitly accepted that risk. There's a Jira ticket, an approval chain, documentation.

Without that context, someone on the team spends an hour re-investigating something that's already been resolved. Multiply that by hundreds of findings, across multiple cloud accounts, across multiple teams, and you start to understand why security teams feel like they're running on a treadmill.

> The teams winning in 2026 will be those who've solved not just the data collection problem, but the data interpretation challenge that comes with it. Visibility without context is just noise.

Chapter 5

Where AI Can Actually Help

Teams want copilots, not autopilots. The three capabilities that matter most, and why practical automation beats moonshot promises.

Where AI Can Actually Help

Despite the adoption challenges, security teams are clear-eyed and pragmatic about where AI can deliver value right now. Rather than chasing moonshots, they're focused on automating the foundational tasks that absorb enormous amounts of analyst time.

The Top Priorities

Capability	Priority
Threat detection and prioritization	57%
Incident triage and response	45%
Compliance reporting and audits	44%
Knowledge sharing and analyst copilots	31%
Policy enforcement	29%
Automated remediation	22%
Intelligent triage	21%
Unified visibility	17%
Cross-cloud correlation	14%

These aren't the flashiest use cases, but they're the ones that matter most to exhausted security teams managing thousands of incidents annually. Automating threat detection, triage, and compliance workflows can reduce noise, speed up investigative cycles, and allow analysts to focus on higher-value, context-rich work that demands human judgment.

What AI-Powered Triage Actually Looks Like

To understand why these capabilities matter so much, consider what intelligent triage means in practice. Today, when a critical finding surfaces — say, an unencrypted EBS volume attached to a production EC2 instance — an analyst has to manually determine: Is this instance internet-facing? What data does it process? Is there a compensating control like application-layer encryption? Is this a known exception with an approved risk acceptance? Has this finding appeared before, and what did the team do last time?

An AI-powered triage system can resolve these questions in seconds by correlating the finding against network topology data, data classification tags, existing exception records, and historical remediation patterns. Instead of presenting a raw finding, it presents a verdict: "This EBS volume is attached to an instance in a private subnet processing non-sensitive batch workloads. The instance has no public IP, no internet gateway route, and is only accessible via a VPC endpoint from the data pipeline account. Previous occurrences were accepted with a 90-day review cycle. Last reviewed January 3rd. Risk acceptance is current." That's the difference between a 45-minute investigation and a 10-second read.

For compliance reporting, the impact is equally tangible. AI can continuously map cloud resource configurations against framework requirements — automatically generating evidence that all production databases enforce TLS 1.2+, that CloudTrail is enabled across all regions with log file validation, that S3 bucket policies deny unencrypted uploads, or that IAM password policies meet CIS Benchmark standards. What used to take a team days of manual evidence collection before an audit can be generated on demand.

What Teams Actually Want: AI Copilots, Not Autopilots

When asked what specific capability they want most from agentic AI, 26% prioritized AI copilots for analysts — tools that reduce noise, accelerate decisions, and unburden overstretched teams. DevSecOps professionals were 29% more likely to prioritize this capability.

This is a critical signal. Teams aren't asking for AI to replace them. They're asking for AI to eliminate the toil — the context-gathering, re-triage, hunting for institutional knowledge — so their best engineers can spend time on judgment calls, not data assembly. A team of five with the right AI support can operate with the leverage of a team of fifteen.

> Most teams aren't looking to rebuild their security stack from scratch — they're looking for relief. The promise isn't revolution; it's sustainable, incremental improvement that compounds over time.

PROWLER

Chapter 6

What Will Separate Leaders from Laggards

The dividing line isn't sophistication — it's pragmatism. What non-adopters need to see, and the emerging playbook for 2026.

What Will Separate Leaders from Laggards

The survey data paints a clear picture of what will differentiate the teams that pull ahead in 2026 from those that fall further behind. The dividing line isn't sophistication — it's pragmatism.

What Non-Adopters Need to See

Among organizations that haven't adopted agentic AI, the barriers to entry are remarkably consistent:

67%

Cost Savings

The deciding factor for adoption

47%

Faster Remediation

Need proven speed improvements

37%

Compliance Gains

Require measurable posture improvement

These aren't abstract goals. They're the metrics that CISOs present to boards and CFOs when justifying security investments. Among current holdouts, just 11% report being very likely to adopt agentic AI within the next year — meaning the burden of proof falls squarely on vendors to demonstrate measurable, defensible ROI.

The Financial Services Signal

Financial services stands out as a leading indicator. Teams in this sector were 22% more likely than average to have already deployed agentic AI into cloud operations. Given the industry's regulatory complexity and the direct financial impact of security failures, this early adoption pattern suggests that high-stakes environments create the urgency needed to move past the trust barrier.

The Emerging Playbook

Based on the survey findings, the teams that will lead in 2026 share a common approach. They prioritize reclaiming analyst time through targeted automation of triage, context gathering, and compliance documentation. They invest in real-time visibility that goes beyond data collection to deliver contextual interpretation. They build institutional memory into their tools, so organizational knowledge doesn't evaporate when people change roles. And they choose solutions that integrate with existing ecosystems rather than demanding wholesale migration.

The teams that fall behind will be those waiting for perfect conditions, perfect tools, or perfect trust before acting. The survey makes clear that the gap between leaders and laggards is already wide and growing wider.

PROWLER

Chapter 7

The Case for AI-Native Cloud Security

From scanner to teammate: what AI-native security must deliver, and why open-source foundations create the flywheel effect.

The Case for AI-Native Cloud Security

Every security vendor will put “AI” on their website this year. That’s the reality of the market. But the survey data suggests that security teams can see through surface-level claims. They’re looking for AI that solves their actual problems, not AI that generates marketing copy.

From Scanner to Teammate

The evolution security teams need is clear: from tools that generate findings to systems that interpret them. The journey goes from “here are your 2,000 findings” to “here are the three things that actually matter in your environment right now, here’s why, here’s what we’ve done about similar issues before, and here’s the recommended action.” That’s the difference between a scanner and a teammate.

What Matters in an AI Security Platform

Contextual intelligence. The system needs to understand your environment specifically — not just flag generic best-practice violations. It should know when a finding has already been reviewed, what compensating controls exist, and whether a risk has been formally accepted. It should understand that an open security group on port 443 means something very different on a public-facing ALB than on an internal microservice in a private subnet. It should correlate IAM policy changes with CloudTrail activity to distinguish between a legitimate infrastructure-as-code deployment and a potential privilege escalation. That level of environmental awareness eliminates hours of redundant investigation every day.

Institutional memory. AI that learns from your organization’s history transforms how teams operate. When a vulnerability recurs, the system should know exactly what the team did last time — that the Log4Shell remediation in your Java services required coordinating with three application teams, that the dependency update broke the health check endpoint in staging, and that the rollout took four days because of a downstream Kafka consumer compatibility issue. Or it flags that it’s never seen this before, which is equally valuable because it means there’s no existing playbook and the team should create one.

Attack path awareness. Rather than presenting findings as a flat list of individual misconfigurations, the system should map how vulnerabilities, misconfigurations, and excessive permissions chain together into exploitable paths. It should identify that the combination of a public-facing workload, an overprivileged service account, cross-account role assumption capabilities, and access to secrets in Parameter Store creates a path to lateral movement — even if no single finding in that chain is rated critical. This is where AI-driven graph analysis and attack simulation fundamentally change prioritization.

Transparency and trust. With 46% of teams doubting autonomous AI, the trust question isn’t optional. In security-critical environments, black-box AI is a non-starter. Platforms built on open-source foundations give teams the ability to audit checks, validate logic, run assessments in their own environments, and verify

AI-driven recommendations against real telemetry. That structural transparency isn't a feature — it's a competitive advantage.

Community-driven accuracy. AI recommendations grounded in checks and guidance that thousands of practitioners have reviewed, challenged, and improved in production are fundamentally more reliable than AI trained purely on documentation and best-practice guides. Battle-tested knowledge is different from theoretical knowledge.

The Flywheel Effect

The most powerful AI security systems will be those that create a flywheel: practitioner knowledge makes the AI smarter, and the AI's insights flow back to improve the tools practitioners use. This virtuous cycle is only possible when the foundation is open, community-driven, and grounded in real-world operational data — not locked behind proprietary walls.

Methodology

This research was conducted by Prowler in partnership with Kickstand, surveying 633 cybersecurity professionals in December 2025 and early 2026. The survey was conducted at 95% confidence with a ±4% margin of error.

Respondent Profile

Geography: 32% United States, 32% United Kingdom, with additional respondents from France, Germany, Italy, Spain, Australia, New Zealand, and Brazil.

Company size: 38% from mid-sized companies (100–1,000 employees), with representation across all size segments from under 100 to over 5,000.

Industries: SaaS (35%), Financial Services (20%), Healthcare, and other sectors.

Roles: Security Analysts (38%), Security Engineers (25%), DevSecOps (13%), CISOs (10%), and other cybersecurity roles.

Qualifiers: All respondents were employed full-time in cybersecurity roles and spoke fluent English.

About Prowler

Prowler is the world's most widely adopted open-source cloud security platform, automating security and compliance across modern cloud environments. With over 13,000 GitHub stars and 45 million downloads, Prowler helps teams reduce triage time, continuously generate audit-ready compliance evidence, and apply AI-driven context — all while maintaining full visibility and control of their environments.

Prowler Cloud delivers zero infrastructure management, SSO and granular RBAC, unlimited historical trends and compliance reporting, and unlimited LighthouseAI usage for insights, triage, and remediation guidance. Deploy in minutes and achieve faster ROI. By combining open-source transparency with enterprise-grade automation, Prowler enables security teams to scale impact without scaling headcount.

Learn more: prowler.com
Try Prowler Cloud: cloud.prowler.com
See Prowler in Action: prowler.com/interactive-demo