Qualys

Cloud Security Forecast 2026

What 2025 Exposed, What 2026 Will Amplify, and How Security Teams Must Adapt

ABOUT THIS REPORT

The Cloud Security Forecast 2026 examines how structural changes in cloud architecture, identity management, automation, and artificial intelligence are reshaping enterprise risk. Drawing on observed incident patterns, market data, regulatory signals, and cloud security maturity research, this report explains why cloud security failures in 2025 were not anomalies but indicators of deeper systemic weaknesses that will be amplified in 2026.

This report is written for security leaders, risk executives, and technology decision-makers responsible for governing cloud environments at scale. Its purpose is not to catalog tools or vulnerabilities, but to **explain how cloud risk materializes, why traditional security decision models fail under modern conditions, and which operating model shifts are required to remain effective.**

Table of Contents

Executive Brief: The Cloud Security Inflection Point

Cloud security has entered a phase of structural acceleration. The defining challenge for organizations is no longer whether cloud risk is understood in principle, but whether it can be governed at the speed, scale, and complexity at which modern cloud environments operate.

In 2025, a series of high-profile cloud incidents across industries demonstrated a consistent pattern, emerging alongside the growing use of automation and AI-assisted tooling by both attackers and defenders. Attackers did not rely on novel vulnerabilities, advanced malware, or technically sophisticated exploits. Instead, they repeatedly succeeded by leveraging AI automation to abuse legitimate access paths, delegate trust relationships, expose credentials, and delay response mechanisms that were already present in cloud environments.

These incidents were not remarkable for their novelty. They were remarkable for their reliability: the same identity and trust failure produced consistent compromise across different environments.

In 2026, the same access paths remain viable. What changes is the velocity and scale at which they can be discovered, evaluated, and exploited. Advances in automation and the adoption of agentic AI systems compress cloud attack timelines from days or weeks into minutes. Security programs optimized for visibility, documentation, or periodic review find themselves structurally misaligned with this reality.

Three foundational forces now define cloud security outcomes:

1. **Identity has replaced the network as the primary control plane**
2. **The rate of cloud change has exceeded the limits of human governance**
3. **Third-party and SaaS trust relationships have expanded the effective attack surface beyond organizational boundaries**

Together, these forces create a condition where risk is driven less by missing controls and more but by decision latency. In 2026, many cloud breaches will occur not because organizations failed to detect risk, but because they could not reduce exposure fast enough.

The Accelerating Cloud Security Challenges

Why Cloud Security Is Evolving Rapidly

Cloud environments do not behave like traditional infrastructure. They behave more like living systems, continuously changing through automation, orchestration, and software-defined control planes. In these environments, security outcomes are determined less by static hardening and more by how identity, permissions, and APIs are governed over time.

Three Structural Realities Now Dominate Cloud Security

Consistently Governing Identity Controls

In cloud-native environments, valid identities, human, workload, or third-party, provide more power than exploits. Once authenticated, attackers operate through legitimate APIs, where actions are valid by design and often indistinguishable from normal administration.

While major cloud providers such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) continue to introduce increasingly granular identity controls including workload identities, pod-level identities, and fine-grained access policies, compromises persist. The challenge is no longer the absence of identity controls, but the difficulty of governing them consistently across cloud, SaaS, and hybrid environments.

As identity models expand to include ephemeral workloads, federated access, and on-premises integrations, misconfigurations and over-privileged relationships continue to provide attackers with durable entry points.

Cloud Change-Rate Outpaces Human Governance

Infrastructure-as-code, ephemeral workloads, and continuous deployment introduce constant drift that degrades even well-designed policies. Permissions, network paths, and service configurations evolve daily or even hourly, often faster than governance processes can validate or enforce them.

Although cloud platforms provide mechanisms to define guardrails, the pace of operational change increasingly exceeds the capacity of human-driven controls, creating windows of exposure that are difficult to detect and even harder to close manually.

Third-Party and SaaS Dependencies Expand the Attack Surface

SaaS platforms, OAuth integrations, and third-party services operate outside hardened infrastructure while retaining deep access to data, workflows, and control planes. These integrations are often trusted by design, with persistent permissions granted to support business operations.

As organizations expand their SaaS footprint, the number of implicit trust relationships grows. When these relationships are misconfigured or compromised, attackers bypass core defenses entirely and gain disproportionate access through legitimate integration paths that are rarely monitored with the same rigor as internal systems.

Why 2026 Is a Key Year to Watch

2026 marks a convergence point where several structural shifts reach maturity at the same time. Artificial intelligence moves beyond assistive tools into agentic systems that execute workflows autonomously. These systems function as identities that authenticate, reason, and act across cloud environments, exposing the limits of security models designed for human-paced decision-making.

At the same time, compliance regimes evolve from static documentation to enforceable operational obligations, while security investment concentrates on cloud and AI risk reduction despite cost pressure. Regulation increases the cost of slow response, automation accelerates both attack and defense, and budgets favor platforms that shorten exposure time rather than add alerts.

What 2025 Signaled and What It Didn't

The most significant cloud security incidents of 2025 exposed a clear pattern: compromise occurs when security decisions are fragmented, manual, or disconnected from identity context.

Across industries and cloud platforms, attackers succeeded by exploiting:

- **Authenticated identities with excessive permissions**
- **Over-privileged workloads, including containers and service accounts, operating with implicit trust**
- **Long-lived credentials and exposed secrets**
- **Delegated access through OAuth and SaaS integrations**
- **Delayed response caused by manual investigation workflows**

These incidents demonstrated that existing identity, workload, and network access paths are already sufficient for compromise. Containers, ephemeral workloads, and service-to-service communication expanded the blast radius once initial access was obtained, while permissive network connectivity enabled lateral movement without triggering traditional alerts.

Importantly, 2025 did not reveal a surge in novel cloud vulnerabilities or advanced exploitation techniques. Public incidents instead confirmed that misaligned identity permissions, workload trust, and network exposure combined with slow response are enough to reliably achieve impact. Entering 2026, the challenge is not awareness. It's execution.

What 2025 Demonstrated and Why It Matters in 2026

2025 SIGNAL	WHAT IT DEMONSTRATED	WHY IT MATTERS IN 2026
Identity-Based Access Dominated Incidents	Authenticated identities bypass preventive controls	Agentic systems can discover and reuse identities continuously
SaaS and OAuth Abuse Enabled Large-Scale Data Theft	Trust relationships define blast radius	Automated exploitation of integrations scales impact
Credential Exposure Led to Disproportionate Access	Secrets often equal control-plane authority	Agents can scan and validate credentials persistently
Privilege Escalation Relied on IAM Relationships	Escalation is configuration-driven	Graph-based reasoning favors automation
Detection Existed But Response Lagged	Speed determined outcome	Agentic attacks compress kill chains

The weaknesses exposed in 2025 align directly with the operating strengths of agentic systems, where reconnaissance, identity mapping, and permission analysis become continuous and autonomous. 2025 was a proof point. 2026 determines how far this model scales.

Key Cloud Security Findings From 2025

Qualys' 2025 cloud security maturity research, conducted across 250+ organizations, indicates that cloud adoption is outpacing the operating models required to govern identity, permissions, and response, creating significant risk as organizations entered 2026.

KEY FINDINGS 1 Cloud Adoption Has Outpaced Cloud Security Maturity

WHAT THE DATA SHOWS

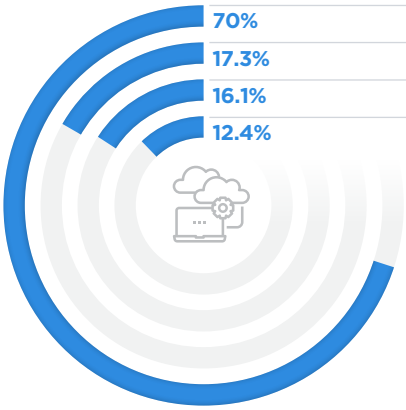

Over **70%** of organizations actively use multi-cloud platforms, Kubernetes, APIs, and cloud-native services

17.3% have implemented Cloud Infrastructure Entitlement Management (CIEM)

16.1% use attack-path analysis

12.4% apply unified risk scoring

WHY THIS MATTERS

Cloud environments are increasingly dynamic and identity-driven, while security models remain static and asset-centric. This mismatch becomes more consequential as automation accelerates in 2026.

KEY FINDINGS 2 Supply Chain and Zero-Day Malware Risks Are Increasing

WHAT THE DATA SHOWS

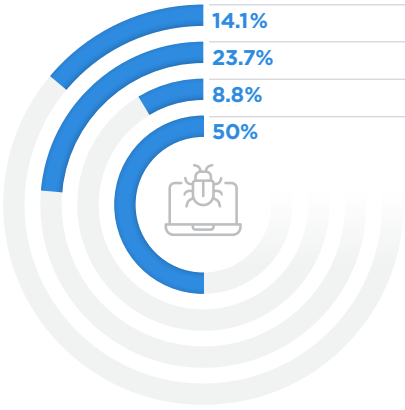

14.1% of organizations have secrets, SBOM, or software composition analysis integratec into CI/CD pipelines

23.7% scan container registries and build pipelines before deploym

8.8% report having runtime visibility or protection capable of detecting unknow or zero-day behavior

Nearly **50%** still rely on manual investigatior and response after alerts are generated

WHY THIS MATTERS

These gaps allow malicious or compromised components such as weaponized open-source packages in ecosystems like NPM to enter cloud environments through trusted build pipelines, where preventive controls often fail. As demonstrated by multiple NPM supply-chain incidents in 2025, zero-day malware increasingly bypasses static checks, making runtime detection and rapid response critical entering 2026.

KEY FINDINGS 3 Human-in-the-Loop Operations Remain the Norm

WHAT THE DATA SHOWS

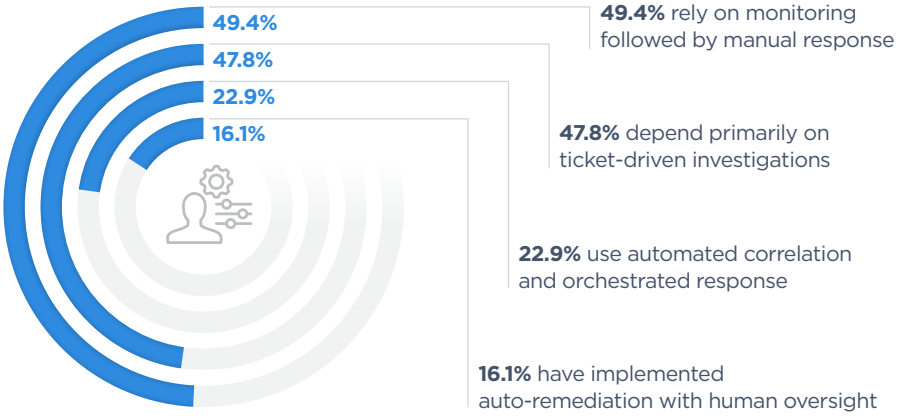

WHY THIS MATTERS

In many organizations, human-in-the-loop operations are driven by compliance requirements and trust considerations, not lack of tooling. However, as attackers automate execution, response models that require manual review introduce security latency debt, causing remediation to occur after impact rather than before it.

KEY FINDINGS 4 Risk Prioritization Remains Severity-Led, Not Blast-Radius-Led

WHAT THE DATA SHOWS

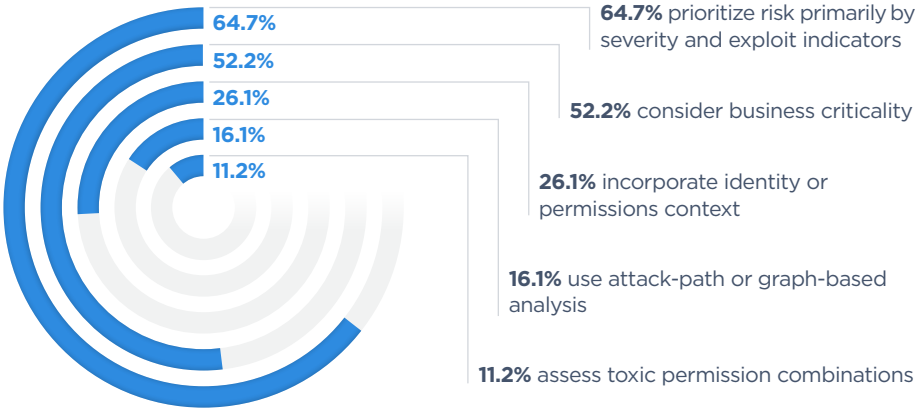

WHY THIS MATTERS

Most organizations cannot evaluate how individual issues combine into viable attack paths. Agentic attackers capable of graph-based reasoning consistently outperform severity-led triage models.

KEY FINDINGS 5 Shift-Left Controls Exist, but Enforcement Is Weak

WHAT THE DATA SHOWS

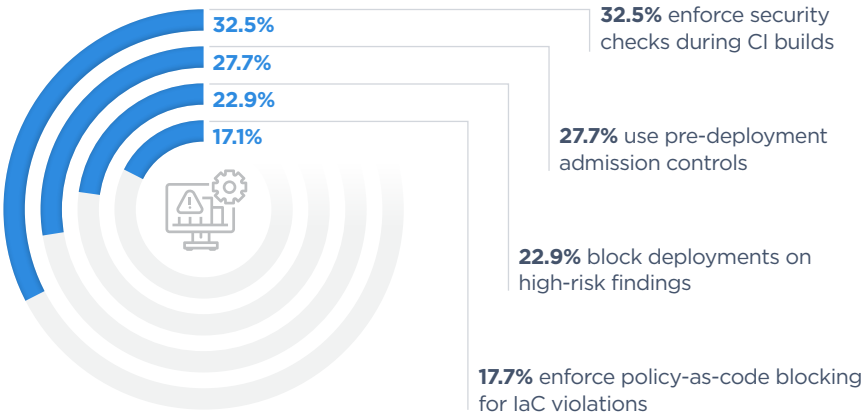

WHY THIS MATTERS

Detection without enforcement creates a false sense of control. In automated cloud environments, attackers exploit what pipelines permit, not what they merely flag.

KEY FINDINGS 6 AI and LLM Security Maturity Is Critically Low

WHAT THE DATA SHOWS

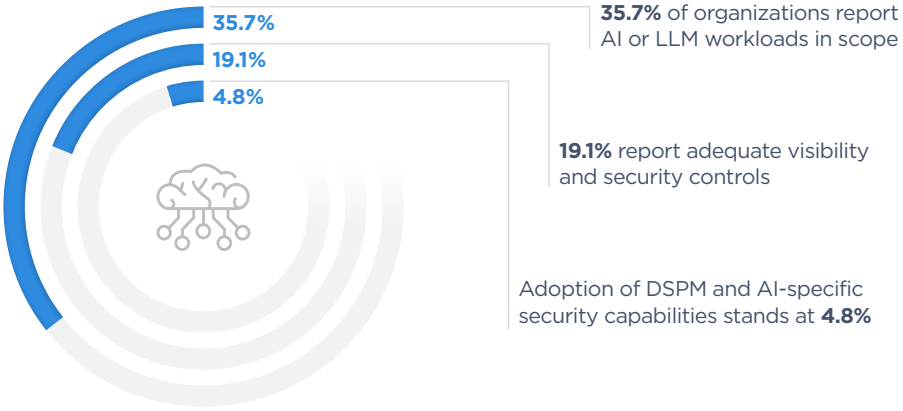

WHY THIS MATTERS

AI systems introduce new identities, data access paths, and trust assumptions. Without governance and visibility, these environments become high-value targets for agentic exploitation.

Synthesis Signal: Security Latency Debt

Taken together, these findings show that cloud risk is being created at machine speed while decisions and remediation remain largely human-paced. With nearly half of organizations relying on manual response, less than 17% automating remediation, and identity complexity expanding rapidly, organizations accumulate a gap between exposure creation and exposure reduction. This gap compounds over time into predictable windows of exploitation.

This is where exposure-centric risk management becomes decisive. Qualys Enterprise TruRisk™ Management (ETM) is designed to collapse security latency debt by continuously prioritizing risk through identity context, attack paths, and business impact. By aligning detection, prioritization, and remediation into a single decision flow, ETM enables security teams to reduce exposure at cloud speed without sacrificing governance.

In 2026, cloud risk will be defined by **identity complexity, automation speed, and decision latency,** not by vulnerability volume.

Market Dynamics and Business Drivers

Cloud security risk is increasingly shaped by forces outside the security organization itself. Budget allocation patterns, regulatory pressure, geopolitical conditions, and competitive expectations now influence which security strategies are viable in practice. Entering 2026, these forces are unusually aligned in their impact on cloud security decision-making.

KEY FINDINGS 1 Budget Signals: What Boards Are Choosing to Fund

Security investment priorities are shifting despite sustained cost pressure, with AI security and cloud security emerging as top investment areas over the next twelve months.

PwC's 2026 Global Digital Trust Insights highlights this trend clearly. In its reporting, **AI security is the top cyber investment priority (46%),** followed by **cloud security (33%)** over the next twelve months. The release shows a consistent pattern, with AI security ranked first and cloud security close behind.

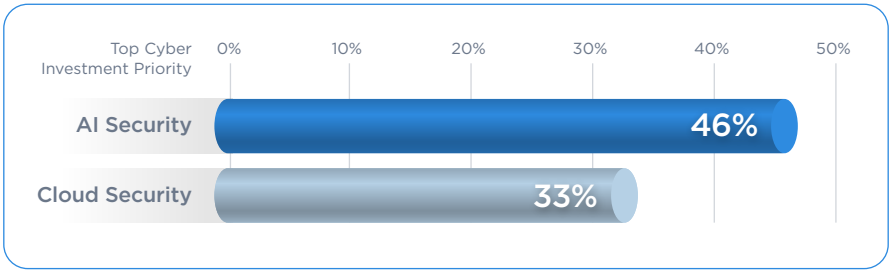

The importance of this signal lies less in the percentages and more in the prioritization logic behind them. Boards are directing spend toward capabilities that reduce systemic business risk across data protection, operational continuity, regulatory exposure, and customer trust. As expectations shift, investment is moving away from tool counts and coverage metrics toward measurable reductions in exposure time, blast radius, and recovery effort, favoring approaches that demonstrate clear operational risk reduction.

KEY FINDINGS 2 Market Growth Signals: Direction Over Precision

Public market forecasts consistently project strong, double-digit growth in cloud security spending across multiple segments. While estimates vary in absolute size and timeframe, they converge on several structural drivers: continued cloud adoption, identity-centric threat models, regulatory pressure, and platform consolidation.

Published Estimates Illustrate This Trend:

The **Cloud Network Security** market is estimated at **USD 13.65 billion in 2025**, forecast to reach **USD 49.86 billion by 2030**, reflecting a **29.58% of CAGR**.

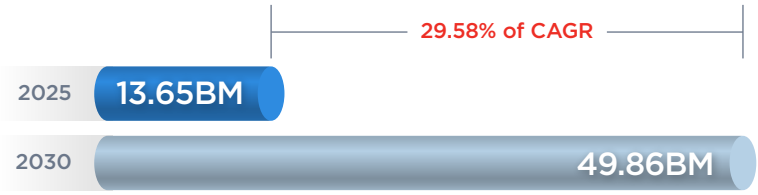

CNAPP market estimates show similar momentum, with one forecast projecting growth from **USD 15.0 billion in 2025 to USD 51.2 billion by 2032** at a **19.2% of CAGR**.

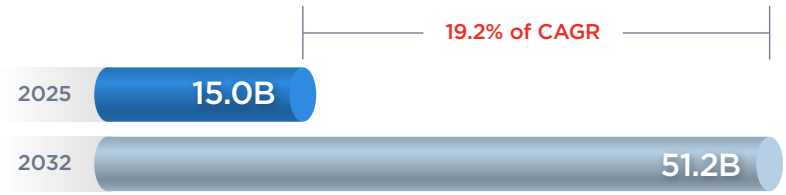

Another CNAPP forecast estimates growth from **USD 10.29 billion in 2024 to USD 65.63 billion by 2034**, representing approximately **20.36% of CAGR**.

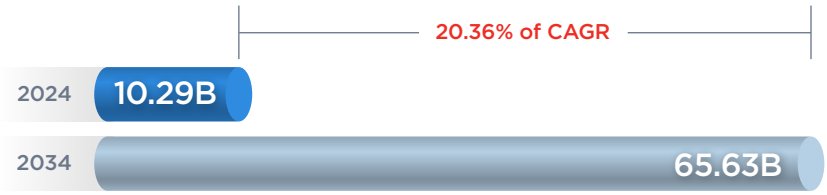

The significance of these figures lies not in their variance but in their agreement on direction. Cloud security is a rapidly growing category driven by the adoption of cloud services, identity-centric threat models, expanding compliance requirements, and increasing demand for platform consolidation.

KEY FINDINGS 3 Geopolitical Drift and Architectural Fragmentation

By 2026, cloud security risk will be shaped less by isolated geopolitical events and more by geopolitical drift. As governments increasingly treat cloud infrastructure as strategic assets, organizations are forced to distribute workloads across regions and providers to meet data sovereignty, regulatory, and national risk requirements.

This drift introduces risk not through new technologies, but through architectural fragmentation. Identity systems are duplicated across regions. Trust relationships expand to local providers and SaaS platforms. Centralized visibility and enforcement erode as environments diverge.

Each architectural decision is rational in isolation. Together, they increase identity and trust complexity in ways that are difficult to reason manually. In this environment, access paths multiply faster than security teams can govern them, amplifying the effectiveness of automated and agentic attacks.

KEY INSIGHT

In 2026, geopolitical forces will influence cloud security not through conflict, but through architectural drift that expands identity and trust complexity.

KEY FINDINGS 4 Economic and Competitive Pressure

Economic pressure continues to compel security leaders to justify investments through measurable business outcomes, including reduced breach probability, decreased downtime, faster recovery, and lower audit and regulatory risk.

At the same time, cloud security is increasingly becoming a competitive differentiator. Customers, partners, and regulators treat security posture as a signal of trust and operational maturity. This aligns with PwC's digital trust framing, where security investment is tied to resilience and confidence rather than technical compliance alone.

As a result, cloud security decisions are increasingly evaluated not by coverage metrics, but by their ability to demonstrate control, continuity, and trust at scale.

WHAT THIS MEANS FOR 2026

Taken together, these drivers narrow the range of viable cloud security operating models. In 2026, security programs are increasingly judged on their ability to reduce exposure time, enforce controls across fragmented environments, govern identity and third-party trust at scale, and support regulatory requirements with operational evidence rather than static reporting.

The Threat Landscape Transformation: Vectors, Actors, and Velocity

Cloud threats entering 2026 are not defined by new classes of attack, but by changes in **speed, scale, and execution model**. The techniques used by attackers are largely familiar. What has changed is how efficiently they are applied and how quickly they converge on impact.

This transformation is driven by automation, identity-centric architectures, and the increasing use of agentic systems by adversaries. Together, these forces compress the cloud attack lifecycle and expose the limitations of security programs built around human-paced investigation and response.

Dominant Trends Shaping Cloud Security in 2026

Analysis of incident patterns and enterprise adoption curves points to four dominant trends shaping cloud security in 2026.

TREND 1 Identity-First Security Becomes Non-Negotiable

Organizations are increasingly focused on mapping and governing identities across cloud environments, including:

- Human identities (SSO, MFA, device trust)
- Workload identities (Kubernetes, serverless, ephemeral services)
- Service accounts and service principals
- Third-party OAuth applications and SaaS integrations

As identity replaces the perimeter, access governance becomes the primary control plane for cloud security. Preventing compromise increasingly depends on understanding who and what can act—not just what is exposed.

TREND 2 Consolidation Around Platforms Accelerates

Point tools can surface issues, but they rarely resolve them in isolation. Platform strategies—such as CNAPP-style approaches—continue to gain traction because they correlate identity context, workload behavior, misconfiguration, data sensitivity, and attack paths into a single decision framework.

As cloud environments grow more interconnected, organizations increasingly prioritize platforms that can reason across signals rather than manage alerts in silos.

TREND 3 Continuous Control Replaces Periodic Posture

Security effectiveness is increasingly judged by how quickly organizations reduce exposure time—from misconfiguration, credential exposure, or malicious code introduction to detection and remediation. Periodic assessments and static posture checks are insufficient in environments where change is constant.

This shift is reinforced by supply-chain incidents observed in 2025, including malicious NPM packages and malware-tainted container images, which entered environments through trusted pipelines rather than external attack vectors. In 2026, continuous control—spanning CI/CD, runtime, and response—becomes essential as prevention alone cannot stop zero-day and dependency-based threats.

TREND 4 Agentic AI Redefines Risk Operations

Security operations are evolving beyond manual alert triage toward AI-driven risk operations that continuously evaluate exposure, identity relationships, and attack paths. Rather than reacting to alert volume, teams increasingly direct agentic systems to prioritize risk, validate exploitability, and initiate containment or remediation actions under human oversight.

As attackers adopt agentic techniques to automate reconnaissance and execution, defensive models must evolve from reactive SOC workflows to proactive Risk Operations Centers (ROC). In 2026, effective programs will use agentic AI not only for detection, but for continuous risk reasoning, exposure reduction, and protection against emerging threats—operating at machine speed while preserving governance, trust, and compliance controls.

Observed Attack Patterns Entering 2026

Qualys threat research from 2025 shows that cloud compromises rarely depended on novel vulnerabilities. Instead, attackers consistently achieved impact by exploiting access paths that already existed inside enterprise environments. Valid identities, permissive configurations, and trusted integrations provided sufficient leverage without advanced exploitation.

These patterns confirm that existing identity, workload, and trust relationships are already sufficient to enable cloud compromise. As automation accelerates in 2026, these same paths will scale faster, increasing impact unless organizations reduce exposure time and govern access more effectively.

How Cloud Attack Vectors Are Changing

Attack vectors in 2026 share four defining characteristics:

Faster

Agentic automation compresses reconnaissance, exploitation, and exfiltration into significantly shorter windows.

Quieter

Attackers increasingly favor API-driven actions over noisy malware-based techniques.

More Identity-Centric

Tokens, OAuth applications, federation abuse, and service identities replace host-level compromise.

More Supply-Chain-Driven

Compromising a single integration or vendor can provide access to multiple downstream environments.

A consistent theme in forward-looking threat analysis is that adversaries are fully embracing AI to automate steps across the attack lifecycle.

Agentic Attacker Adoption and Focus Areas

Agentic attackers operate as autonomous systems capable of continuously evaluating, prioritizing, and executing actions across cloud and SaaS environments. Rather than relying on single exploits, these systems optimize for return on access, selecting identity and configuration driven paths that provide durable control with minimal noise. Recent threat research and incident analysis show that these techniques are already being used manually and semi automatically; agentic execution simply removes human pacing from the loop.

The focus areas below consistently deliver the highest return on investment for attackers and align closely with observed real-world incidents.

1 Reconnaissance

MITRE ATT&CK (Cloud): Discovery (TA0007)

Agentic attackers begin with continuous, API driven reconnaissance across cloud, SaaS, and identity planes. Rather than scanning hosts, they enumerate:

- Cloud accounts, projects, and subscriptions
- IAM roles, service accounts, and trust relationships
- Publicly exposed storage, APIs, and endpoints
- OAuth applications, consented scopes, and token behaviors

This mirrors patterns observed in multiple cloud breaches where attackers first mapped identity and service relationships before taking action. Qualys cloud threat research and third-party incident analyses repeatedly show that attackers rely on **misconfigurations and exposed access paths**, not zero day vulnerabilities, to identify viable targets.

The agentic advantage is persistence: reconnaissance is continuous, not episodic. Agents re-evaluate environments as configurations change, allowing them to exploit newly introduced exposure within minutes of creation.

2 Privilege Escalation

MITRE ATT&CK (Cloud): Privilege Escalation (TA0004), Lateral Movement (TA0008)

Rather than exploiting software flaws, agentic attackers construct privilege graphs from IAM policies, role inheritance, federation rules, and OAuth scopes. These graphs reveal escalation paths created by:

- Pass-role or assume-role permissions
- Token exchange and workload identity federation
- Overlapping SaaS and cloud admin privileges
- Delegated access through CI/CD pipelines and automation accounts

This technique aligns with real-world incidents in which attackers escalated privileges using only valid permissions, often chaining together low-severity IAM misconfigurations. Qualys and other industry research highlight that privilege escalation in cloud environments is typically **configuration-driven**, not exploit-driven.

Agentic systems excel here because graph reasoning scales better than human reviews. They can evaluate thousands of potential paths and select the lowest-friction route to impact.

3 Persistence

MITRE ATT&CK (Cloud): Persistence (TA0003)

Once access is achieved, agentic attackers prioritize persistence through configuration rather than malware. Common techniques include:

- Creating or modifying service accounts and managed identities
- Generating long-lived access keys, refresh tokens, or OAuth grants
- Embedding access into CI/CD workflows or automation scripts
- Leveraging identity federation that survives credential rotation

These methods have been observed in multiple cloud incidents where attackers retained access even after partial remediation. Because persistence is achieved through legitimate configuration changes, it often survives incident response actions focused on user credentials or endpoint artifacts.

Agentic attackers continuously validate persistence paths and re-establish access when disrupted, increasing defender fatigue and response cost.

4 Data Discovery and Exfiltration

MITRE ATT&CK (Cloud): Collection (TA0009), Exfiltration (TA0010)

With durable access in place, agentic attackers shift to targeted data discovery. Rather than bulk extraction, they focus on:

- High value storage buckets, data lakes, and object stores
- SaaS platforms containing customer, financial, or identity data
- Analytics services and AI/ML data pipelines

Exfiltration is performed through legitimate APIs, often throttled or segmented to avoid triggering volume-based alerts. This behavior aligns with several high-profile SaaS and cloud data breaches in 2024-2025, where attackers used approved integrations and export mechanisms rather than custom tooling.

Qualys and third-party cloud breach analyses consistently show that **API-native exfiltration** reduces detection and attribution, especially in environments where behavioral baselines are weak.

Geopolitical Risk and Threat Actor Focus

Geopolitics affects cloud security in three concrete ways.

Sovereignty and Data Residency Pressure

Organizations adopt regional cloud controls, restrict administrative access by location, and implement tighter audit requirements.

Supply-Chain Fragility

Sanctions, regulatory divergence, and vendor dependencies increase systemic risk. Compromise of a single SaaS integration or widely used service can ripple across multiple organizations.

Threat Actor Focus Shifts

Nation-state and criminal operations increasingly exploit geopolitical tension for espionage, disruption, and influence, a pattern consistently highlighted in forward-looking threat reporting.

Pace Comparison: Governments, Enterprises, and Attackers

A defining dynamic entering 2026 is the growing mismatch between attacker speed and defender decision cycles.

- Governments and regulators operate on legislative timelines measured in months or years.
- Enterprises operate on quarterly delivery and planning cycles.
- Attackers, increasingly agentic, operate on minute-to-hour cycles.

By the time teams review findings, approve remediation, and coordinate response, attackers may have already achieved persistence or data access.

In cloud environments governed by APIs and identity, speed determines outcome more reliably than control coverage.

IMPLICATIONS FOR SECURITY TEAMS

The threat landscape in 2026 does not reward incremental improvements to alerting or visibility. It rewards operating models that reduce decision latency, reason over identity and trust relationships, and execute response actions at machine speed under human oversight.

The next sections examine how architectural choices and agentic systems amplify these dynamics, and what security teams must change to remain effective.

Multi-Cloud and Zero Trust: Foundational Shifts in Architecture

By 2026, multi-cloud and Zero Trust are no longer strategic aspirations. They are structural conditions that determine whether cloud risk is governed or absorbed as unmanaged exposure.

The Security Implications of Multi-Cloud Reality

Multi-cloud adoption continues to accelerate, driven by resilience, regulatory requirements, cost optimization, and placement of AI and data-intensive workloads. While these drivers are rational, multi-cloud increases security risk through inconsistency rather than platform count.

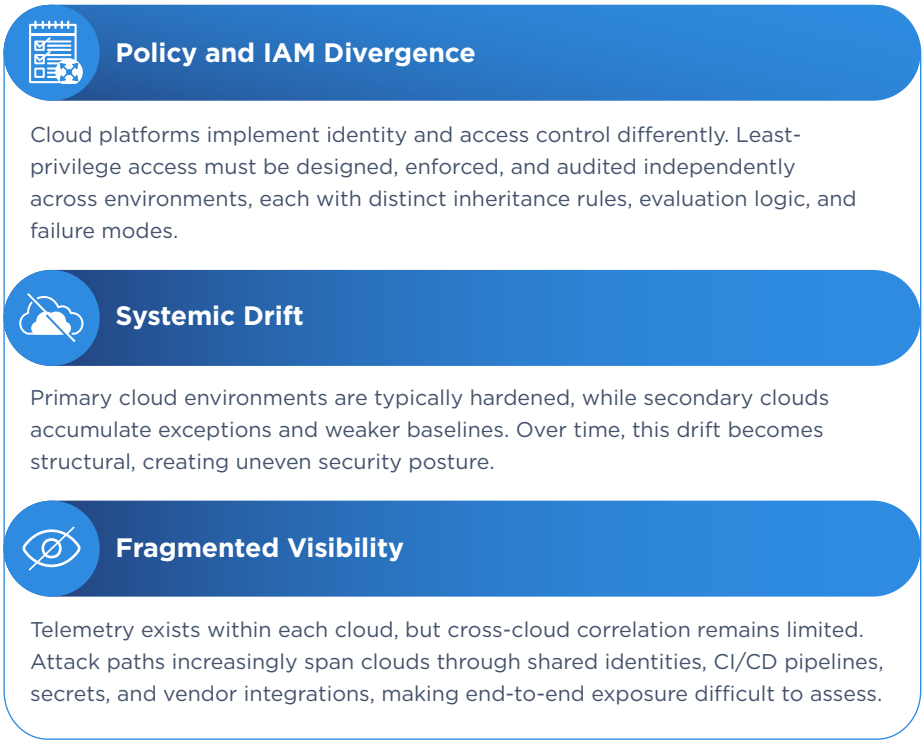

What Changes in 2026

Multi-cloud security maturity becomes a visible differentiator. Auditors, regulators, and customers increasingly expect consistent control enforcement across all environments.

Zero Trust in the 2026 Landscape

In 2026, Zero Trust becomes an operational requirement rather than an architectural ideal. Effective implementation depends on continuous identity verification, enforced least privilege, just-in-time administrative access, context-aware decisions, and identity-based segmentation.

Regulatory pressure further reduces tolerance for static trust models and persistent access. As a result, Zero Trust is defined by operational behavior, not diagrams, with dynamic trust adjustment and rapid revocation becoming core capabilities.

Where Multi-Cloud and Zero Trust Break Down

Most organizations struggle to operationalize Zero Trust consistently across multi-cloud environments. User access is prioritized, while non-human and workload identities remain statically governed, and security harmonization is deferred in favor of availability and cost.

This leads to identity sprawl, uneven enforcement, and fragmented visibility. Attackers exploit these gaps rather than defeating Zero Trust itself. By 2026, the challenge is not adoption, but consistent governance of identity and trust at scale.

ARCHITECTURAL IMPLICATIONS IN 2026

Multi-cloud and Zero Trust amplify each other's risks when implemented inconsistently but reinforce resilience when governed coherently. Security teams must treat identity, access, and trust relationships as shared control planes across environments, not as provider-specific implementations.

Agentic AI in Defense and Attack: The Security Imperative

In 2026, cloud security is no longer defined by the strength of individual controls, but by the speed and coherence of decision-making across increasingly automated environments. As cloud platforms become more identity-driven, API-operated, and interconnected, both attack and defense shift from episodic actions to continuous execution.

Agentic AI accelerates this shift. By removing human pacing from core security operations, it changes where cloud compromise originates, how quickly it scales, and what effective defense requires. The following forecasts outline how agentic systems reshape cloud security dynamics and where their impact will be most concentrated.

FORECAST 1 Agentic AI is the Force Multiplier Behind Cloud Compromise

Agentic AI becomes the dominant force shaping both cloud defense and cloud compromise. As AI systems evolve from rule-based automation to autonomous agents capable of reasoning and adaptation, security outcomes shift from isolated control effectiveness to continuous decision-making at machine speed.

1. **Decision Logic Shifts from Playbooks to Identity-Centric Optimization.**

   Agentic systems do not execute predefined workflows. They evaluate available actions in real time based on likelihood of success, cost, and impact. In cloud environments, this decision logic consistently favors identity-centric paths. IAM graphs, non-human identities, control-plane APIs, CI/CD pipelines, and SaaS integrations offer high leverage with low friction. These systems define authority, persist over time, and operate through legitimate interfaces, making them ideal substrates for agentic execution.

2. **Agentic Offense Scales by Operating Inside Trusted Control Paths.**

   On the attackers' side, agentic adversaries function as autonomous red teams. They continuously map IAM relationships to identify privilege escalation paths, monitor non-human identities for excessive or unused access, probe cloud APIs to test permission boundaries, and exploit CI/CD workflows and OAuth integrations to achieve durable access. Rather than breaking controls, they operate within them, selecting the lowest-cost path that blends into normal operations.

3. **Agentic Defense Becomes the Only Viable Counterweight to Automation at Scale.**

   On the defenders' side, agentic AI acts as a force multiplier for security teams facing identity complexity and automation scale. Autonomous defenders correlate signals across clouds, reason over identity and permission graphs, assess business impact, and propose containment and remediation actions in near real time. This accelerates SOC transformation by shifting analysts from manual triage to supervision of automated validation, prioritization, and response.

The significance of 2026 is economic. Agentic AI collapses the cost of reconnaissance, iteration, and persistence while increasing the speed at which access paths can be discovered and reused. As cloud environments grow more fragmented and identity-driven, the balance shifts toward actors that can reason, adapt, and execute continuously. Security effectiveness increasingly depends on whether organizations can apply agentic capabilities to defense as rapidly and coherently as attackers apply them to exploitation.

2 Where Agentic Attacks Will Concentrate

By 2026, cloud compromise will concentrate less on infrastructure flaws and more on systems that define authority, automation, and trust at scale.

1. Identity and access management becomes the primary attack surface.

In cloud environments, power is encoded in IAM relationships, not hosts. Permissions form large, interconnected graphs that change continuously and are rarely reviewed holistically. Agentic systems excel here, reasoning across roles, trust paths, and inheritance to identify low-cost escalation routes that are invisible when permissions are assessed in isolation. Seemingly benign rights combine into toxic access paths, enabling privilege escalation and data exfiltration through legitimate APIs.

2. Non-human identities emerge as a dominant target class.

Service accounts, workload identities, CI/CD roles, and managed identities increasingly mediate cloud access, often without MFA, with broad permissions, and limited monitoring. Once compromised, these identities provide durable, low-noise access that bypasses user-centric detection models. Agentic attackers can discover, exploit, and maintain these identities continuously, making persistence easier than remediation.

3. Control planes, pipelines, and integrations amplify impact.

Cloud APIs, CI/CD workflows, and OAuth-based SaaS integrations offer predictable, authorized mechanisms to create resources, change permissions, and extract data. Agentic attackers abuse these paths not by breaking controls, but by operating within them, modifying automation, harvesting secrets, and extracting data at scale. By 2026, durable cloud compromise increasingly occurs through trusted workflows and integrations rather than overt intrusion.

IMPLICATIONS FOR CLOUD SECURITY TEAMS

Agentic attackers do not need to break into cloud environments. They move through access paths that already exist, faster than human decision cycles allow.

Effective defense in 2026 requires continuous reasoning over identity and permissions, strict governance of non-human identities, behavioral monitoring of API usage, and first-class treatment of CI/CD and SaaS as attack surfaces.

Most importantly, security teams must shift from human-in-the-loop execution to human-on-the-loop oversight. Manual operations cannot scale against autonomous threats.

The next section grounds these dynamics in real-world incident archetypes and examines how organizations can address the residual risk left behind by 2025 while preparing for what comes next.

Incident Archetypes and Fix Patterns: What 2025 Left Behind

The most instructive cloud incidents observed in 2025 were not outliers. They followed repeatable archetypes that expose where modern cloud security programs remain structurally weak. These incidents matter not because they were novel, but because they reveal which failure modes persist even in mature environments.

Understanding these archetypes allows security teams to address residual risk quickly while preparing for the acceleration expected in 2026.

ARCHETYPE 1 Third-Party System and Social Engineering

OBSERVED PATTERN

In several high-impact incidents, attackers gained access through third-party cloud systems or vendor-managed workflows. Social engineering was used to manipulate identity recovery, change approval processes, or data export actions. Once inside, attackers operated with legitimate access, resulting in broad customer impact.

WHAT FAILED

Vendor and partner identities often hold elevated privileges but fall outside core security monitoring. Recovery and administrative workflows are frequently optimized for availability and support efficiency rather than adversarial resilience. When these processes are compromised, traditional infrastructure controls offer little protection.

FIX PATTERN

Effective mitigation focuses on constraining trust rather than expanding detection:

- Enforce least-privilege access for vendor identities
- Isolate third-party platforms from sensitive data stores
- Require multi-party approval for high-impact actions such as exports and identity recovery
- Monitor vendor identities for anomalous access patterns

These controls reduce blast radius even when initial access is obtained.

ARCHETYPE 2 OAuth and SaaS Integration Abuse

OBSERVED PATTERN

Attackers compromised OAuth tokens associated with third-party SaaS integrations, enabling large-scale data theft across customer environments. The infrastructure itself remained intact. The attack operated entirely through approved integrations.

WHAT FAILED

OAuth applications are often granted broad scopes during setup and rarely reviewed afterward. Tokens persist indefinitely, and refresh behavior is poorly monitored. Once compromised, these integrations provide durable access that bypasses traditional perimeter and endpoint controls.

FIX PATTERN

Mitigation requires treating SaaS and OAuth as first-class attack surfaces:

- Inventory all OAuth applications and remove unused integrations
- Enforce scoped permissions and short token lifetimes
- Apply conditional access controls for bulk actions such as mass export or query
- Monitor token refresh and query behavior for anomalies
- Detect anomalies in query behavior, token refresh activity, and data egress patterns

This pattern reinforces a critical shift: business SaaS and integrations must be governed with the same rigor as core cloud infrastructure.

Clearing 2025 Exposure Debt While Preventing 2026 Backlog

Many organizations enter 2026 carrying unresolved exposure from prior years. These "leftover" issues persist not because they are unknown, but because remediation cannot keep pace with change.

Common Sources of Recurring Risk include:

- **Excessive privileges and permission creep**
- **Unmanaged accounts, projects, or subscriptions**
- **Inconsistent logging and monitoring coverage**
- **Ungoverned third-party OAuth integrations**
- **Weak identity recovery and change control processes**
- **Backlog-driven remediation where findings grow faster than fixes**

Preparing for 2026 Requires a Dual-Track Approach.

1. **Reduce Existing Exposure Quickly**

   Organizations must aggressively shrink exposure windows by cleaning up excess access, decommissioning unused resources, and closing high-impact attack paths left behind from 2025.

2. **Prevent the Next Backlog**

   Automation and governance must be embedded so that new risks do not accumulate faster than they can be addressed. Without this shift, every improvement simply delays the next saturation point.

The objective is not perfect prevention. It is sustained reduction in exposure time as environments continue to accelerate.

IMPLICATIONS FOR CLOUD SECURITY TEAMS

The incidents of 2025 demonstrate that cloud security failures are rarely caused by lack of tools or awareness. They result from operating models that allow risk to persist long enough to be exploited.

Organizations that treat these failures as isolated events will repeat them. Those that treat them as systemic signals can materially reduce risk entering 2026.

Strategic Imperatives and Security Team Priorities for 2026

The cloud security challenges in 2026 are not primarily technical. They are operational. The patterns observed in 2025 and the acceleration driven by automation and agentic systems make clear that incremental improvements to tooling or coverage will not be sufficient.

Security teams must adopt a small number of decisive shifts in how risk is modeled, prioritized, and reduced.

The Five Priority Shifts Security Teams Must Make

IMPERATIVE	WHAT MUST CHANGE	WHY THIS MATTERS	HOW TEAMS CAN ACT
Shift From Asset-Centric to Identity-Centric Security	Infrastructure can no longer be treated as the primary unit of risk. In cloud environments, identities—human, service, workload, and SaaS—form the true control plane.	While cloud adoption is widespread, only one in four organizations incorporate identity context into risk prioritization. In identity-driven attack models, this leaves escalation paths invisible until they are exploited.	• Continuously inventory all identities and their permissions • Monitor permission drift and toxic combinations • Prioritize risk based on who can reach what, not asset severity alone
Move From Alert-Driven Response to Attack-Path Awareness	Alert volume and severity-based triage are no longer sufficient. Teams must understand how individual issues combine into viable attack paths.	Incidents in 2025 showed that attackers rarely exploit single findings in isolation. They chain misconfigurations, credentials, and trust relationships to reach impact.	• Model attack paths across cloud, SaaS, and identity layers • Focus remediation on breakpoints that collapse entire paths • Measure success by reduced exposure, not reduced alert count
Replace Human-in-the-Loop Bottlenecks with Human-on-the-Loop Oversight	Manual review should not gate every security action. Humans must supervise automated decisions rather than execute each step.	A significant portion of organizations still rely on manual response. Against agentic attackers, this introduces delays that are consistently exploited.	• Automate low-risk, high-confidence remediation • Reserve human decision-making for ambiguous or business-critical actions • Tune automation based on outcomes, not alert volume
Enforce Security Decisions Earlier and With Authority	Detection without enforcement is ineffective. Security teams must be empowered to prevent high-risk configurations from reaching production.	Although many organizations perform CI/CD checks, only a minority of these enforce blocking. This gap creates exploitable windows of exposure.	• Shift enforcement left using policy-as-code • Block deployments that introduce dangerous identity or exposure paths • Align DevOps incentives with secure outcomes, not deployment speed alone
Prepare for Agentic Defenders to Match Agentic Attackers	Manual operations cannot scale against autonomous threats. Defensive automation must evolve beyond scripts into goal-driven systems.	Agentic attackers excel at continuous reconnaissance, reasoning, and execution. Static defenses fall behind even when coverage appears strong.	• Invest in systems that reason over identity, permissions, and behavior • Use automation to prioritize and execute risk reduction continuously • Treat AI governance as a security control, not merely a compliance requirement

THE CLOSING PERSPECTIVE

The defining challenge of cloud security in 2026 is not awareness. It is operating at the speed required to govern identity, trust, and automation effectively.

Agentic attackers will not break in. They will move through access paths that already exist, faster than human decision cycles allow. Security teams that adapt their operating models accordingly will reduce risk materially. Those that do not will continue to detect incidents after impact.

How Cloud Security Programs Succeed in 2026

The cloud security failures observed in 2025 did not occur because organizations lacked visibility, tooling, or intent. They occurred because security operating models were built for a slower, more linear threat environment.

Cloud environments are now governed by identity, automation, and APIs. Attackers have adapted accordingly. They no longer rely on breaching hardened perimeters or exploiting novel vulnerabilities. They move through legitimate access paths, abuse delegated trust, and operate at speeds that exceed human decision cycles.

The acceleration expected in 2026 is not speculative. It is the predictable outcome of converging forces: agentic automation, expanding identity complexity, architectural fragmentation, and regulatory pressure that demands faster, more provable control. In this environment, security latency becomes the dominant risk factor.

Reducing this latency requires more than incremental improvement. It requires a reset in how cloud risk is modeled and managed. Asset inventories must give way to identity graphs. Alert queues must give way to attack-path reasoning. Manual execution must give way to supervised automation.

This shift does not remove risk. It compresses exposure time, constrains blast radius, and reestablishes control in environments that outpace human tracking. Without it, organizations will continue to detect compromise after impact, even as their security stacks expand.

In 2026, cloud security effectiveness will be defined less by what organizations can see, and more by how quickly they can decide and act.

Insights and statistics presented in this report are derived from responses to the Qualys 2026 Cloud and Application Security Maturity Survey conducted in collaboration with 250+ global organizations and reflect the operational experiences and practices reported by participating enterprises.

Talk to a Qualys expert to get started on priority driven cloud risk management with Qualys TotalCloud™

Book a Demo

Qualys.

About Qualys

Qualys, Inc. (NASDAQ: QLYS) is a pioneer and leading provider of disruptive cloud-based security, compliance and IT solutions with more than 10,000 subscription customers worldwide, including a majority of the Forbes Global 100 and Fortune 100. Qualys helps organizations streamline and automate their security and compliance solutions onto a single platform for greater agility, better business outcomes, and substantial cost savings. For more information, please visit qualys.com. Qualys, Qualys VMDR® and the Qualys logo are proprietary trademarks of Qualys, Inc. All other products or names may be trademarks of their respective companies.