Futurum®

Google Cloud

Closing the AI Confidence Gap:

Cloud-Native Security as a Key to Agentic AI Adoption

APRIL 2026

The Executive Conflict: Innovation vs. Execution

AI has captured the industry's attention faster than any prior technology trend. It brings lofty productivity expectations, to immense AI improvement and implementation efforts, to changes of cybersecurity's very nature.

As organizations rapidly adopt AI, Chief Information Security Officers (CISOs) and their teams are increasingly focused on ensuring that security evolves in parallel. Rather than treating security as a downstream checkpoint, leading organizations are integrating AI security practices alongside deployment, aligning innovation with governance to support scalable, responsible adoption.

The instantiation of all this AI adoption activity is primarily cloud-based. While some enterprises are building their own "AI Factories," the use of AI capabilities from cloud hyperscalers including the roll-out of autonomous agents remains a key architectural choice. Every large cloud hyperscaler continues to invest heavily in AI capacity, as AI has definitively crossed the threshold from isolated experimentation to widespread enterprise rollouts.

Figure 1. Dual Mandate to Balance Aggressive Innovation with Profound Risk Mitigation

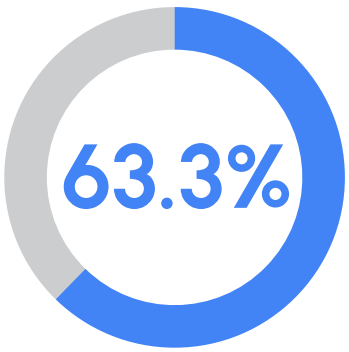

of CIOs prioritize investing in emerging technologies (AI, ML, IoT) as their top strategic initiative – the most consistently dominant priority across all of CY25.

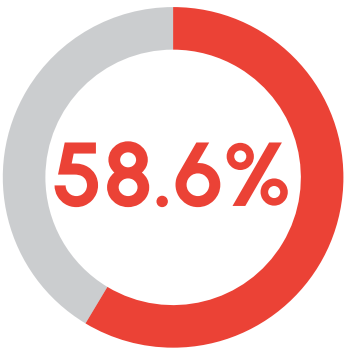

of CIOs cite data security, information leaks, and privacy risks as their top AI concern, even as cybersecurity normalizes into an operational baseline.

Source: Futurum Research, CIO Insights Global Survey (Q4 2025)

The enterprise is currently navigating a structural reset from deploying early, experimental generative AI tools to operationalizing Agentic AI. Unlike basic assistive applications, agentic workloads are designed to plan, act, and autonomously orchestrate core business workflows and transactions. As these agents integrate directly into systems of record (often at scale), the tolerance for error collapses, requiring strict controls and governance: traditional role-based access, forensic logging, and zero-trust principles, adapted to agentic needs of scale, speed, and complexity.

Despite this momentum, a critical "confidence gap" persists: a Cloud Security Alliance survey found that 73% of organizations lack confidence in their ability to execute an AI security strategy[1]. Security teams must rapidly evolve their operational maturity around securing AI; rather than acting as bottlenecks, they must collaborate closely with AI developers to accelerate time-to-production and ultimately time-to-value.

Practitioners must expand their expertise across at least three distinct fronts:

- **Master the AI Stack:** Security teams must understand and tailor governance to the nuances of AI infrastructure, data, models, platforms, and agents. Even before AI, security flaws often emerge from subtle differences in how systems actually behave versus how people think they behave. AI doesn't change that.
- **Secure the AI Stack:** Security teams must select and deploy necessary controls, practices, and processes throughout the AI lifecycle.
- **Continually Align with Business Use Cases:** Practitioners must understand exactly how their organization intends to use AI. Governance models must accommodate specific operational contexts without stifling innovation; this is an ongoing loop, not a once-and-done effort.

Rather than bolting on third-party security controls, security teams may benefit from working closely with the hyperscaler capabilities around AI security, particularly AI hardware/infrastructure, AI models, lifecycle protection, and unified visibility.

This analysis focuses exclusively on "workload AI," defined as focused on security initiatives aimed at protecting AI applications, proprietary models, and autonomous agents hosted on cloud infrastructure. It does not cover "workforce AI": AI used for cybersecurity or general end-user engagement with workforce AI in SaaS applications.

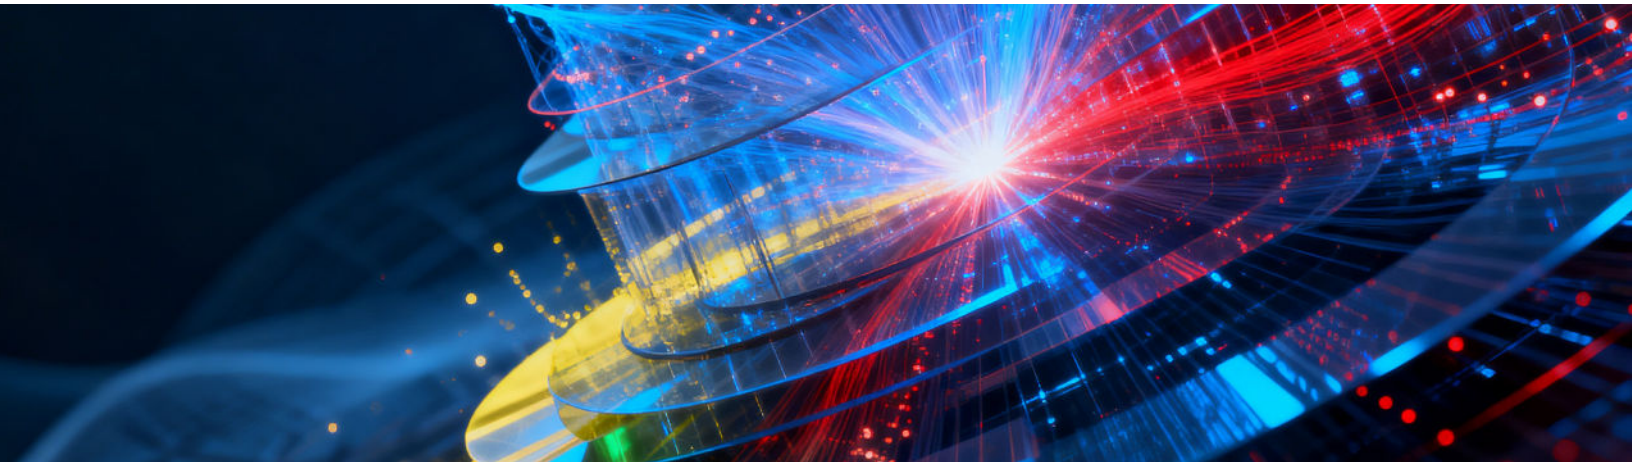

1 Cloud Security Alliance, State of AI Security and Governance Survey Report (2025)

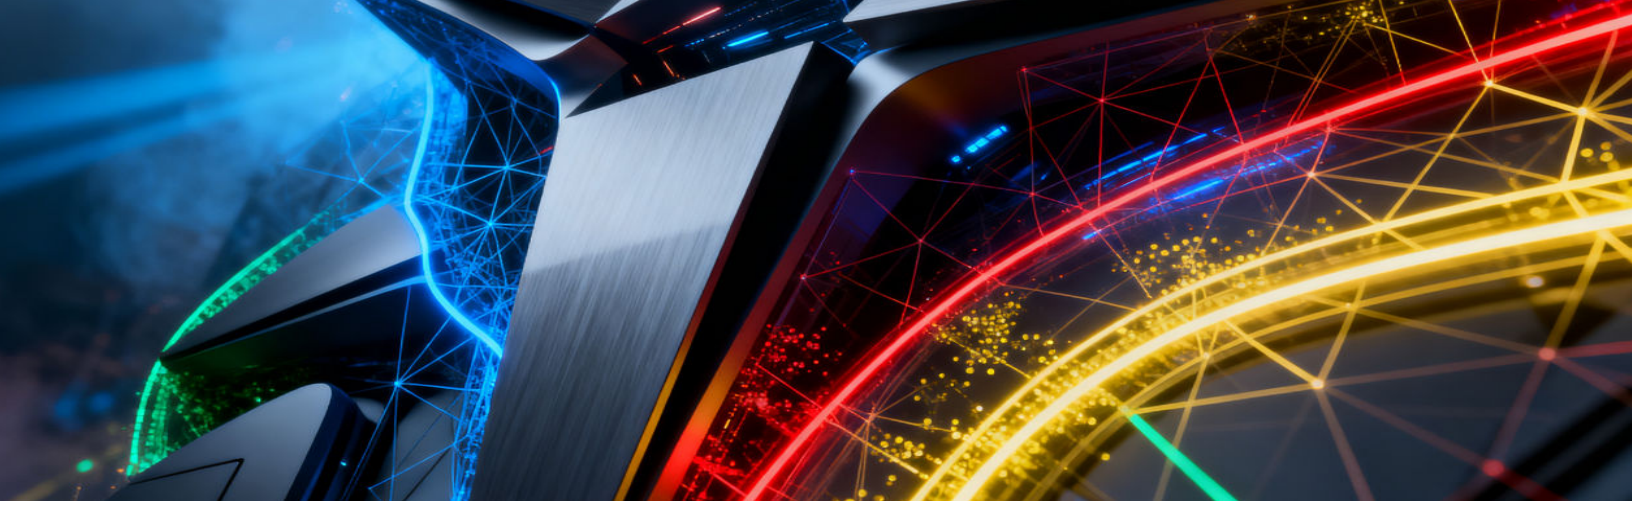

The Challenger View: The "Bolt-On" Paradox

Historically, technology providers delivered IT infrastructure with minimal native security capabilities. Enterprise security teams had to procure and integrate third-party tooling to establish the necessary visibility, access controls, and threat detection capabilities. Over time, enterprise defense added multiple layers of "bolt-on" external tools. While some of this has been necessary to secure hybrid and multi-cloud environments, it doesn't fully address the security needs that are required to be tightly built-in into the cloud fabric.

Modern cloud hyperscalers have fundamentally altered this paradigm. They now engineer deep, integrated security functionality directly into the underlying fabric of their platforms. As enterprises operationalize mission-critical AI workloads, security cannot operate as an external checkpoint. It must be consumed natively.

The Need for an Integrated Approach

Relying on legacy, third-party security architectures alone have some limitations for AI deployments at scale and speed. Traditional security platforms rely on deterministic policy enforcement and fixed controls, which are inherently insufficient for managing the probabilistic and non-deterministic behaviors of generative AI. To secure advanced AI workloads effectively, external multi-cloud solutions must be complemented with cloud-native AI security to overcome three structural challenges:

- **Architectural Blind Spots, Costs, and Latency:** Third-party security tools sit outside the core cloud infrastructure, lacking the internal services visibility needed to sanitize complex AI interactions. In the Agentic AI era, novel vulnerabilities such as prompt injection, tool poisoning, and unintended agent-to-agent escalations occur in-path. Rather than routing inference traffic to a disjointed third-party API—which can disrupt workflow orchestration and introduce latency or transport costs—robust multi-cloud security oversight should be unified with native, in-path cloud protections.

- **Integration Complexity and Fragility:** Piecing together fragmented external security tools with native AI workloads requires balancing traditional APIs with newer components such as Model Context Protocol (MCP) servers and agents. Without a tightly integrated security architecture, this complexity increases the risk that integrations break when the underlying cloud platform updates. A seamless strategy that bridges multi-cloud tools with native platform services reduces this operational fragility.

- **Acquisition Cost and Portfolio Sprawl:** Procuring and managing dozens of highly specialized, disjointed point solutions drastically inflates the Total Cost of Ownership (TCO) and exacerbates technology sprawl. Organizations benefit most from a consolidated approach that pairs a unified multi-cloud security platform with native hyperscaler defenses, drastically reducing enterprise fragmentation.

The Provider Paradox & Trusting the Stack

Continuing to rely on third-party security exposes a systemic contradiction in enterprise cloud strategy: organizations willingly trust cloud hyperscalers to host their most sensitive AI models, proprietary training data, and core applications, yet they hesitate to trust those same platforms to secure their workloads.

Historically, maintaining a physically and logically separate security stack was prudent to ensure air-gapped separation of duties. However, modern cloud operating models and the convergence of enterprise applications with hyperscale infrastructure mean that organizations now benefit more from a streamlined, platform-native security approach.

As organizations transition to persistent, autonomous agents that act directly on systems of record, platform-native defense becomes an architectural necessity. The cloud provider that builds specialized infrastructure, develops platform layers, and hosts model weights is inherently the best-equipped entity to secure the entire AI stack of infrastructure, data, models, platforms, and agents.

Foundational security principles are not abandoned but modernized through automated governance aligned with the cloud-shared responsibility model. Unifying security within the platform reduces integration friction and establishes the control planes needed to safely scale AI, enabling IT, security, and development teams to build and operate secure AI applications and agents using hyperscaler-native tools and services.

The Governance Friction of Legacy Architectures

Applying fragmented, "bolt-on" security to artificial intelligence also creates inefficient governance, leading to disjointed oversight, operational blind spots, delayed time-to-production, and potential downstream liabilities. This inefficiency for traditional governance targets and newer agentic governance requirements becomes hazardous when addressing global regulatory mandates. Multinational enterprises require verifiable data residency, ensuring key aspects of the system remain within strict geographic boundaries to satisfy national sovereignty requirements. Legacy tools outside the core infrastructure generally lack the platform-level control necessary to enforce these jurisdictional boundaries.

The Architecture: Requirements for "Full-Stack" Defense

Enterprise technology leaders are aggressively consolidating platform spend around unified cloud hyperscalers to reduce integration friction. The architectural blueprint for securing artificial intelligence must similarly evolve. Protecting complex, non-deterministic AI systems with legacy tools introduces operational blind spots. To confidently scale agentic AI, organizations should consider implementing a platform-native, "full-stack" security architecture governed by three fundamental technical requirements.

- **Vertical Integration (Chip-to-Agent):** Effective defense requires intrinsic visibility across the entire technology stack, extending from specialized compute infrastructure (TPUs and GPUs) directly up through proprietary model weights, platform services, and agent applications. Modern vulnerabilities do not occur in isolation. The architecture must account for highly complex execution paths; specifically user-to-agent, agent-to-agent, and agent-to-tool interactions. Cloud hyperscalers that seamlessly integrate learnings and specialized guardrails from their own AI research laboratories into their foundational compute infrastructure possess a distinct architectural advantage, particularly against novel, sophisticated threats.

- **Lifecycle Protection (Build, Run, and Use):** Securing AI workloads is not a static deployment checkpoint. It requires a continuous security continuum spanning the build, run, and active use phases. During the build phase, organizations must enforce rigorous software supply chain controls to ensure foundational training sets remain free of poisoned or manipulated data. At run and use time, the architecture demands dynamic runtime protections positioned directly within the inference path for security efficacy and performance.

- **Unified Visibility and Control:** The rapid proliferation of AI deployments renders siloed security operations obsolete. Security teams need centralized, automated capabilities to discover all AI assets across development and production environments. This observability empowers practitioners to visualize data lineage, track model dependencies, and map the expanding attack surface in a single dashboard. Consolidating these signals eliminates the guesswork associated with fragmented tools, providing actionable governance and deterministic oversight for executive leadership confidence.

Strategic Guidance: Sovereignty & The Path Forward

Deploying AI globally introduces acute boardroom concerns regarding regulatory compliance and national sovereignty. For multinational enterprises, particularly in highly regulated regions such as EMEA and JAPAC, operationalizing AI demands verifiable data residency. Executives must demand that their platform provider guarantee both proprietary training data and active model weights remain strictly confined within defined geographic boundaries. A sovereign-ready cloud environment applies hardened access controls to satisfy local jurisdictional requirements without sacrificing computational agility.

However, structural architecture alone is insufficient against motivated adversaries; organizations must also proactively stress-test their models against adversarial attacks prior to deployment. An effective defensive posture necessitates a dual-pronged approach to "Red Teaming." First, organizations should utilize automated red teaming to continuously simulate attack paths and identify toxic combinations. Second, this must be coupled with human-expert-led red teaming to assess architecture, evaluate governance, and uncover sophisticated, novel vulnerabilities through advanced threat modeling.

Strategic recommendations to close the AI confidence gap and safely accelerate autonomous workload deployment:

- **Consolidate to Accelerate:** Move away from fragmented, third-party point solutions deployed specifically to secure AI. Unified, cloud-native security strategy can reduce integration complexity, eliminate architectural blind spots, and offset the acute cybersecurity skills gap.

- **Trust the Stack:** The cloud provider building specialized infrastructure, developing platform layers, and hosting models is inherently best positioned to secure them. Building upon a trusted, vertically integrated stack is the most effective way to embed dynamic guardrails directly into the agentic workflow.

- **Start with Governance:** Verifiable, automated governance from day one unblocks executive hesitation and provides the foundational trust required to confidently scale AI innovation. Begin by: leaning into newer frameworks, such as Google's Secure AI Framework (SAIF); engaging organizations such as the Coalition for Secure AI (CoSAI); and monitoring and supporting emerging AI requirements from established organizations such as ISO and NIST. Then, automate compliance with continuous AI posture controls to enforce secure-by-default configurations and automatically generate evidence.

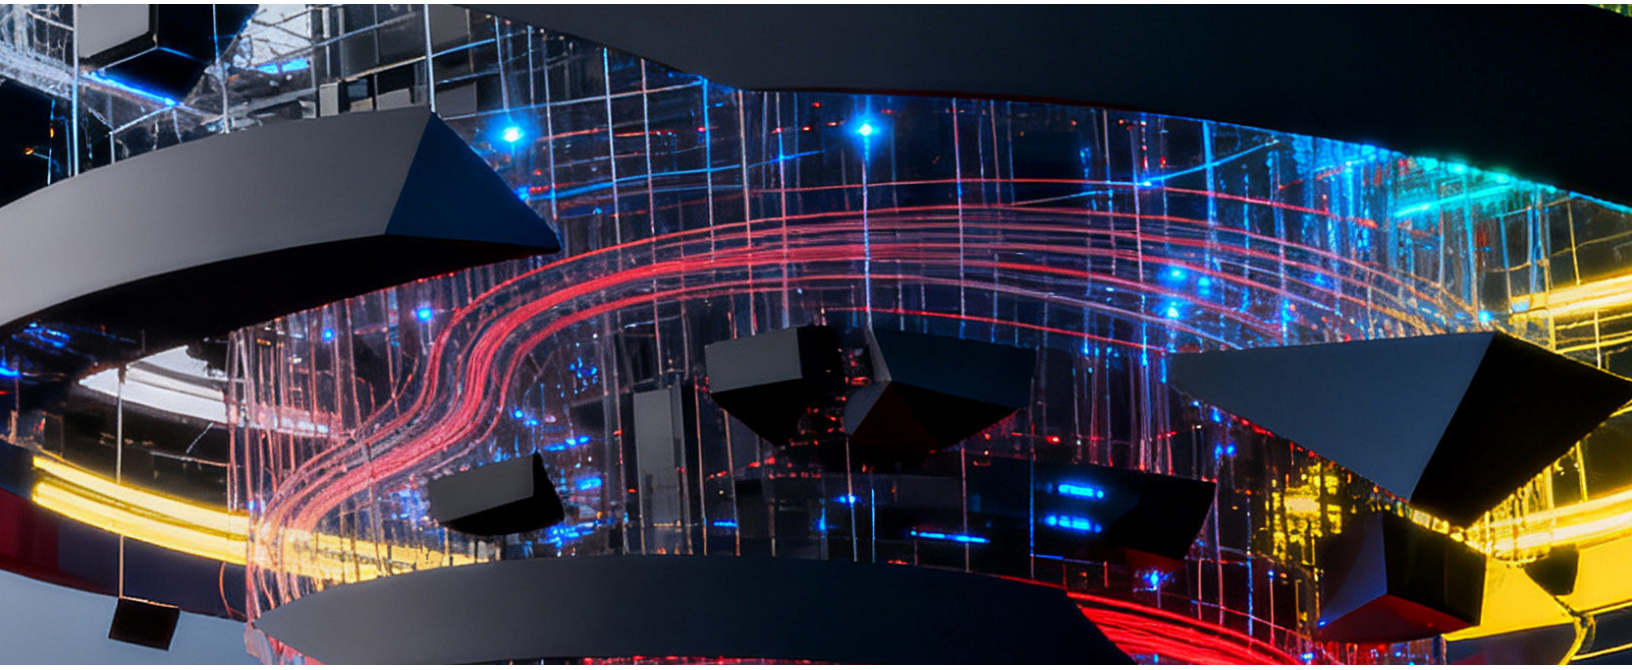

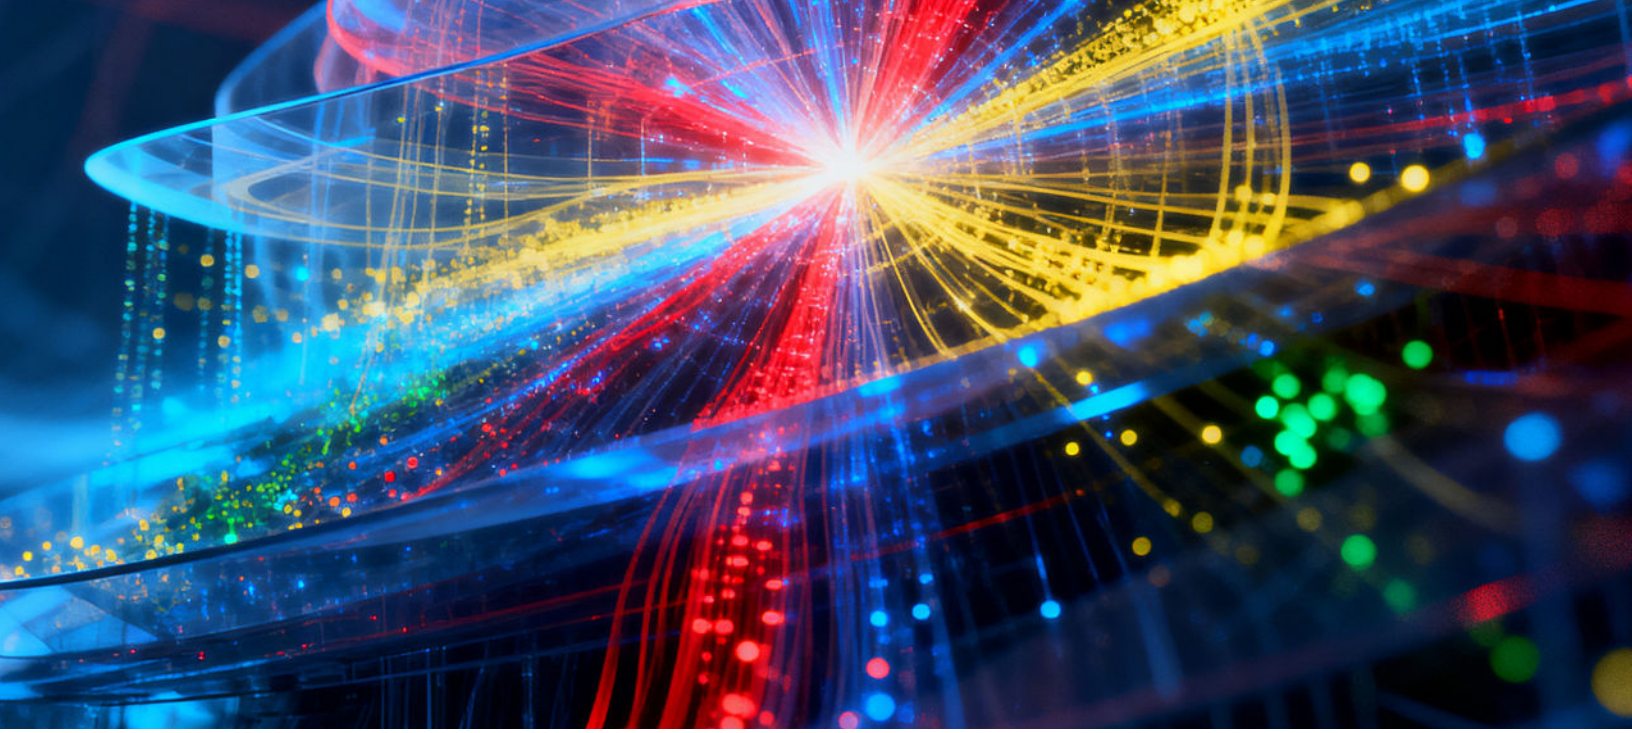

Securing the Full AI Stack with Google Cloud

To effectively address these architectural requirements, Google Cloud offers a comprehensive, platform-native approach that bridges the AI confidence gap. By integrating insights from Google DeepMind, Google Cloud embeds advanced defense mechanisms directly into its foundational infrastructure. This ensures continuous protection across the entire AI lifecycle from build to runtime, and secures the full AI stack, including models, proprietary data, and autonomous agents.

Organizations transitioning to Agentic AI can build securely on services such as Gemini Enterprise and Gemini Enterprise Agent Platform. Google Cloud eliminates the latency and fragility of legacy "bolt-on" tools by delivering unified visibility and automated posture controls natively through Security Command Center and a multi-cloud AI Application Protection Platform (AI-APP) with Wiz. To defend against novel in-path vulnerabilities, Model Armor provides dynamic runtime guardrails that mitigate probabilistic risks such as prompt injection and tool poisoning in real time.

Furthermore, Google Cloud directly supports complex governance and sovereignty mandates through verifiable data residency with Data Boundary and Sensitive Data Protection. When paired with proactive threat modeling and human-led red teaming from Mandiant Consulting, this vertically integrated architecture empowers enterprises to safely manage their expanding attack surface and accelerate AI innovation with complete trust.

Important Information About This Report

AUTHORS

Fernando Montenegro
Vice President & Practice Lead,
Cybersecurity & Resilience | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

Google Cloud

ABOUT GOOGLE CLOUD

Google Cloud is a global cloud platform that enables organizations to build, deploy, and scale applications across infrastructure, data, and AI services. With a strong emphasis on AI innovation, Google Cloud integrates advanced capabilities such as generative AI, machine learning, and data analytics directly into its platform to help enterprises accelerate decision-making and automate workflows. Security is foundational to Google Cloud's architecture, with a zero-trust approach, built-in encryption, and AI-driven threat detection designed to protect data, applications, and users at scale. By combining AI-driven insights with deeply embedded security controls, Google Cloud helps organizations innovate confidently while maintaining resilience in an increasingly complex threat landscape.

Futurum®

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.