

The 90-Day AI Security Reality

Why “Going All-In on AI” Still Comes Back to Identity, Access, DLP, and API Governance

Dr Zero Trust
Chase Cunningham

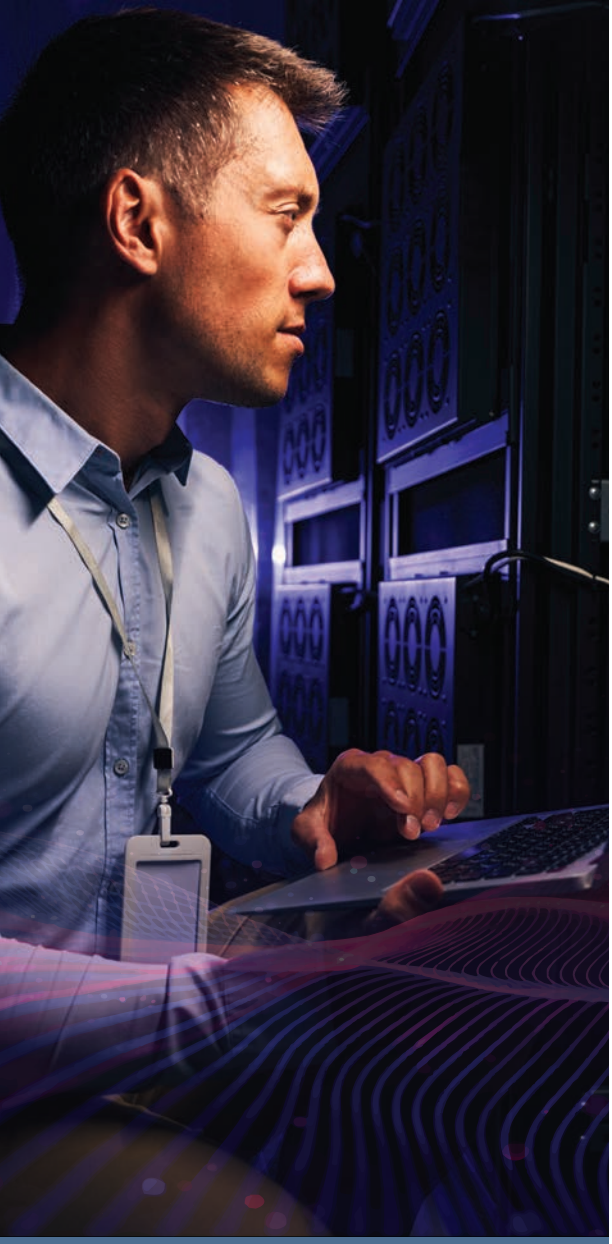
AI Adoption Is Outpacing Security

Executive mandates to accelerate AI adoption are arriving faster than most security teams can operationalize them. Across the enterprise, organizations are being pushed to support shared token budget agreements, internal MCP servers, agentic workflows acting with employee credentials, and sensitive data moving to external models that have not been fully classified or governed. The numbers confirm the squeeze: 78% of technology leaders say AI adoption is outpacing their organization's ability to manage the associated risks, and 85% admit they prioritize speed-to-market over exhaustive AI vetting. Meanwhile, only 7% of organizations worldwide have reached mature AI-security readiness — the lowest-scoring pillar in Cisco's global Cybersecurity Readiness Index.

Only 7% of organizations worldwide have reached mature AI-security readiness



Most security leaders are now being asked to secure an AI-enabled operating model on an aggressive 90-day timeline. The challenge goes beyond "AI risk" to include the rapid expansion of identity sprawl, non-human access, unmanaged API exposure, and uninspected data flows, all under the banner of innovation. Yet the pressure, however new it feels, resolves to familiar fundamentals: enterprises still need to control who can access what, what data can leave, how users and machines interact and under which identities, and where trust boundaries actually exist. In practice, most organizations are layering AI on top of unresolved security basics rather than confronting something fundamentally new.



Enterprises Confront Familiar Gaps Behind New AI Pressure

The evidence is already visible in breach data. IBM's 2025 Cost of a Data Breach Report found that 20% of breached organizations experienced an incident involving "shadow AI" — ungoverned tools operating outside IT oversight — adding an average of $670,000 to breach costs; 97% of organizations that suffered an AI-related incident lacked proper AI access controls, and 63% had no AI governance policy at all. The behavior driving this is pervasive: 72% of enterprise generative AI use is shadow IT conducted through personal accounts, while the volume of data sent to generative AI applications grew more than thirty-fold in a single year. Analysis of real-world enterprise browsing telemetry shows that 40% of files employees upload to generative AI tools contain personal or payment card data, and generative AI has become the single largest channel of corporate-to-personal data exfiltration. These are data classification and DLP failures — familiar problems moving at a new velocity.

Identity tells the same story. Machine identities now outnumber human identities 82 to 1, and 42% of those machine identities carry privileged or sensitive access — yet 88% of organizations still define "privileged user" as applying only to humans. As AI agents join the workforce, they are inheriting exactly the credential patterns security teams have spent a decade trying to retire: 44% of organizations use or plan to use static API keys to authenticate AI agents, 43% rely on username-and-password combinations, and only 18% of security leaders are confident their current identity infrastructure can handle agent identities at all.

Machine identities now outnumber human identities 82 to 1, and 42% of those machine identities carry privileged or sensitive access.

The interfaces carrying all of this traffic are similarly under-governed. Salt Security found that 99% of organizations encountered API security issues in the past twelve months — and, tellingly, 95% of API attacks originated from authenticated sources, meaning compromised credentials and tokens rather than outside intrusion. Roughly 80% of organizations lack continuous, real-time API monitoring, leaving them blind to precisely the machine-to-machine traffic that agents and MCP servers generate.



The current AI wave, in other words, does not necessitate the invention of an entirely new security stack. Even in 2026, the most effective path is to apply proven controls to a new operating context: identity, to know which users, agents, service accounts, and delegated credentials are acting; access, to enforce least privilege, scoped permissions, and policy boundaries for human and non-human actors alike; DLP, to monitor and control what sensitive data is exposed to models, tools, and downstream systems; and API governance, to secure the interfaces through which models, agents, MCP servers, and enterprise systems exchange instructions and data.



How to Secure AI Without Slowing Deployment

Organizations that move fastest on AI without grounding that expansion in these four control planes will create a larger attack surface than they can observe or control. The trajectory is already measurable: 54% of organizations report an experienced or suspected AI-agent security or privacy incident in the past year, yet 85% have no formal accountability structure for agent behavior, and 81% feel pressure to deploy agents even when governance is not fully in place. Requiring every team to solve identity, access, DLP, and API governance independently also slows deployment; preapproved development patterns that build these controls in from the start can instead accelerate the path to production.

"What customers need first isn't a new security model, it's the discipline to treat this as a problem they've solved before," says Alex Vulovic, Principal Solutions Architect at Arctiq, in an interview for this article. "We've run API gateways for 15 years, and the requirements haven't changed: access control, authentication, observability, monitoring." Vulovic says an MCP gateway, such as Cloudflare's, applies an established model to AI traffic, governing how models, agents, MCP servers, and enterprise systems exchange data and instructions. (Arctiq partners with Cloudflare on enterprise Zero Trust and SASE deployments.)

The discipline Vulovic describes must guide which AI deployment requests move forward. As AI moves from experimentation to scaled deployment, tomorrow's winners will not be the ones saying "yes" to every AI deployment request. Successful organizations need rapid adoption while preserving policy enforcement, auditability, and operational trust, treating the 90-day mandate as a forcing function to finally close identity, data, and API gaps that predate AI entirely. This isn't a reason or call to slow down AI, but rather to secure AI deployment using the controls that already underpin resilient enterprise security.

> **"The "AI era" may be new, but the path to managing it still starts with the fundamentals.**

References

1. EY, “EY Survey: Autonomous AI Adoption Surges at Tech Companies as Oversight Falls Behind,” March 4, 2026, https://www.ey.com/en_us/newsroom/2026/03/ey-survey-autonomous-ai-adoption-surges-at-tech-companies-as-oversight-falls-behind.
2. Cisco, “2025 Cybersecurity Readiness Index,” 2025, https://newsroom.cisco.com/c/dam/r/newsroom/en/us/interactive/cybersecurity-readiness-index/2025/documents/2025_Cisco_Cybersecurity_Readiness_Index.pdf.
3. IBM Security, “Cost of a Data Breach Report 2025,” July 30, 2025, https://www.ibm.com/think/x-force/2025-cost-of-a-data-breach-navigating-ai.
4. Netskope Threat Labs, “Cloud and Threat Report: Generative AI 2025,” 2025, https://www.netskope.com/resources/cloud-and-threat-reports/cloud-and-threat-report-generative-ai-2025.
5. LayerX Security, “Enterprise AI and SaaS Data Security Report 2025,” 2025, https://go.layerxsecurity.com/hubfs/LayerX_Enterprise_AI_and_SaaS_Data_Security_Report.pdf.
6. CyberArk, “Machine Identities Outnumber Humans by More Than 80 to 1,” April 23, 2025, https://www.cyberark.com/press/machine-identities-outnumber-humans-by-more-than-80-to-1-new-report-exposes-the-exponential-threats-of-fragmented-identity-security/.
7. Cloud Security Alliance and Strata Identity, “AI Agent Identity Crisis: Standards Emerge as Enterprises Lag,” March 18, 2026, https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/03/CSA_research_note_okta_ai_agent_iam_framework_enterprise_gap_20260318-csa-styled.pdf.
8. Salt Security, “Q1 2025 State of API Security Report,” February 26, 2025, https://salt.security/press-releases/salt-labs-state-of-api-security-report-reveals-99-of-respondents-experienced-api-security-issues-in-past-12-months.
9. Salt Security, “Salt Security Report Shows API Security Blind Spots Could Put AI Agent Deployments at Risk,” October 8, 2025, https://www.prnewswire.com/news-releases/salt-security-report-shows-api-security-blind-spots-could-put-ai-agent-deployments-at-risk-302577909.html.
10. Gravitee, “State of AI Agent Security 2026,” July 2026, https://techcrunch.com/sponsor/gravitee/ai-agents-just-doubled-inside-the-enterprise-confidence-rose-faster-than-control-did/.

About the Analyst: Dr. Chase Cunningham is a globally recognized cybersecurity strategist and creator of the Zero Trust Extended (ZTX) Framework. As former Chief Strategy Officer at Demo-Force.com and Ericom Software, he has led breakthrough innovation and market expansion across the security industry. A trusted advisor to the U.S. Department of Defense and global agencies, Chase shapes cyber policy and national defense strategies. Known as “Dr. Zero Trust,” his insights have redefined how organizations approach security.