

Agents Without Guardrails

The Agentic AI Governance Gap in the Enterprise

August 2026 EMA Research Report

By Christopher M. Steffen, CISSP, CISA, VP of Research
Information Security, Risk and Compliance Management



Table of Contents

1	Introduction
3	Key Findings
5	Voices of the Survey
7	Deployed at Scale, Governed at Pilot
12	Agents are Already Acting Outside Their Intended Scope
18	The Governance Deficit: Why Agents Act Out of Bounds, and What Explains It
25	Conclusion
27	Research Methodologies and Demographics



Introduction

AI agents are everywhere. Enterprises are using them in every facet of their production environments: from resetting passwords, triaging security alerts, generating and committing code, handling customer inquiries, and making autonomous decisions in workflows that touch sensitive systems and data. It is not a question of whether the enterprise will deploy agentic AI at scale, but understanding what happens when those agents do something they were never supposed to do.

Among the enterprise technology and security leaders that took part in this research survey, nearly two-thirds have experienced an AI agent acting outside its intended scope. Nearly one-third report measurable organizational impact – data exposure, operational disruption, financial or reputational consequences. When it happened, most organizations took hours to respond. Many still cannot fully reconstruct what the agent did. Some only found out because a customer told them first.

This is the governance gap exposed by this research: not as a single control failure, but as a compound one. Each of those dimensions has its own data, its own failure pattern, and its own implication for what needs to change.

The organizations managing this well are not the ones that moved fastest. They are the ones that defined what agent behaviors and actions were permitted before deploying them at scale and built the detection and response capability to enforce those boundaries in production.

The data from this research organically comes to three conclusions that EMA will further explore in this report:

1. Agent actions must be governed in real time
2. Detection and containment capabilities must be in place before scaling
3. Decommissioning is a security event, not a project closure task



Key Findings

The following findings represent some of the most interesting findings from the research. They are sequenced to tell the overarching story: adoption is further along than the public narrative suggests, governance has not kept pace, and the consequences of that gap are already being felt in production.

65%

Agents are already doing things they were never authorized to do. Sixty-five percent of enterprises have experienced an AI agent acting outside its intended scope – 29% with measurable organizational impact, 36% a near-miss caught before damage occurred. For a technology most organizations have been running in production for less than two years, this rate should reframe how the industry talks about AI risk.

32%

Most organizations cannot stop it fast – and some learn about it from customers. Only 32% can detect and contain an out-of-scope agent action within minutes through automated mechanisms. Fifty-five percent need hours and manual steps, while 3.5% first learn about rogue agent behavior when a customer or external partner reports it – after the incident has already reached external visibility.

34.2%

Authorization is being checked at the wrong time – or not at all. Only 34.2% evaluate agent authorization at the moment an agent acts. The majority run on standing permissions scoped at provisioning time and are never revalidated for the specific task at hand. This is the direct mechanism through which overprovisioned agents produce unauthorized outcomes.

46%

After an incident, nearly half of organizations cannot reconstruct what the agent did. Just over 46% cannot easily produce a complete audit trail for a specific agent’s last 30 days of activity. Without that record, the questions that matter most after containment – what was accessed, what data was touched, what actions need to be reversed – cannot be answered without significant manual investigation, if at all.

30%

Abandoned pilots are leaving live credentials in production environments. Thirty percent of agentic AI pilots have been paused indefinitely, discontinued, or abandoned. Most were real deployments with real system access and real credentials. Most were not cleaned up. The residual exposure each discontinued pilot leaves behind is ungoverned, unmonitored, and in many cases, unknown.

46%

Deployment has outpaced the infrastructure built to manage it. Forty-six percent of enterprises are already scaling agentic AI across multiple departments. More than 92% report increased AI and bot-driven traffic targeting their APIs. Yet, 47% cannot reliably inventory all deployed agents, 55% lack a fully integrated governance process, and 94% express confidence in a posture the incident data does not support.



Voices of the Survey

Respondents were asked to explain in as much detail as possible their organization’s approach to evaluating and deploying controls around artificial intelligence (AI). Below are some curated responses to this question.

The Hard Reality



CISO, Technology Services

“Our agents work. They do what we deployed them to do, most of the time. What we don’t have is a clean answer for what happens when they don’t – who knows first, how fast we can stop them, what data they touched. That’s the part that keeps me up at night.”

The Accountability Gap



VP of IT, Financial Services

“When something goes sideways with an agent, the first question everyone asks is, ‘whose agent is this?’ And the honest answer, more often than not, is that nobody is completely sure. We have teams deploying agents faster than we can track them, and accountability gets murky fast.”


The Visibility Problem



*Director of Information Security,
Technology Services*

“We can tell you what our agents are supposed to do. We can’t always tell you what they actually did last Tuesday. The audit trail question is the one that makes me nervous – not because we haven’t thought about it, but because when I actually look at what we can reconstruct, it’s less than I’d like.”

The Governance-First Lesson



CIO, Retail & E-Commerce

“We had an agent in a customer-facing workflow take an action we hadn’t anticipated – nothing catastrophic, but it was a wake-up call. We’d been so focused on what the agent could do that we hadn’t been rigorous enough about defining what it should never do. That conversation should have happened before we went to production, not after.”




Deployed at Scale, Governed at Pilot

The enablement framing that dominated enterprise AI conversations for the past two years no longer describes what is actually happening in the enterprises. The overwhelming majority of organizations in this sample have moved past that stage. The challenge they face now is not how to begin deploying agentic AI. It is how to govern deployments that are already running, already touching production systems, and already generating the kind of exposure that governance frameworks were never designed to manage at this speed or scale.


When respondents were asked to describe the current stage of their agentic AI initiatives, 46.0% described themselves as already scaling agentic AI across multiple departments and production workflows. Just over 31% are actively moving prototypes to production. Only 20.8% remain in the research or prototype phase and just 2.5% have paused due to governance or risk concerns.

The concurrent nature of AI adoption compounds the scale. Almost 79% of respondents are running both generative and agentic AI simultaneously, not sequentially. The governance complexity of managing generative AI's outputs and agentic AI's actions is accumulating in parallel, in the same organizations, often managed by the same teams. Agent counts are higher than most organizations publicly discuss. More than 43% of respondents are running between six and 20 distinct active agents, including pilots. Nearly 40% are running more than 20. These are multi-agent portfolios running in core operational functions, and not single-agent proofs of concept.


How would you describe the current stage of your organization's agentic AI initiatives (AI systems that can autonomously plan, take multi-step actions, and access tools or systems with limited human intervention)?



Is your organization currently using generative AI, agentic AI, or both?



What is your best estimate of the total number of distinct AI agents currently active across your organization, including pilots?




The functions enterprises are deploying agents into are not peripheral. IT help desk and service desk automation leads at 56.4%, followed by software development and engineering at 37.6%, security operations at 29.2%, and customer support at 25.2%. An IT help desk agent that resolves tickets autonomously can reset passwords, provision access, modify configurations, and interact with directory services. A security operations agent that triages alerts operates against the same infrastructure that governs the organization’s security posture. These are functions in which an agent acting outside its intended scope does not produce a minor inconvenience. It produces an incident.


The security signal accompanying this deployment is already visible. More than 92% of respondents report an increase in AI and bot-driven traffic targeting their customer-facing applications and APIs. Thirty-six percent describe the increase as significant. The attack surface is not theoretical; it is the current operating environment for virtually every enterprise in this sample, including traffic from authorized agents, misconfigured agents, and adversarial bots indistinguishable from legitimate agent traffic without the right tooling to differentiate them.

Against that deployment backdrop, the governance picture reveals consistent and consequential gaps. Nearly half of respondents cannot reliably tell you what agents are running in their environment. Fifty-three percent have an automated, centralized agent inventory process. Eighteen percent rely on a manual inventory and 28.2% have only partial visibility. An organization that cannot see all of its agents cannot govern them – and 47% of this sample is in that position while simultaneously running dozens of agents in production.


What specific use cases has your organization identified as the highest priority for agentic AI?



Has your organization seen an increase in automated or AI agent-driven traffic targeting your customer-facing applications or APIs in the past 12 months?



How would you describe your organization’s ability to discover and inventory all AI agents currently operating in your environment, including those deployed by individual teams?




Governance integration shows the same pattern. Almost 45% report that AI agent governance is fully integrated into broader IT and security processes. Forty-nine percent report partial integration. Partial integration means exactly what it sounds like: some agents, some environments, and some deployments are governed, while others are not. As business units accelerate agent adoption independently, the ungoverned tail grows faster than the governance perimeter can close.


The partially-integrated population faces the most acute version of the visibility problem. An organization with partial agent inventory and partial governance integration has a structural blind spot that grows proportionally with deployment scale. Every agent deployed by a business unit outside the governance perimeter expands that blind spot without triggering any alert in the systems designed to catch it. The organizations that will face the most difficult remediation exercises in 2026 and 2027 are not the ones that built no governance at all – they are the ones that built partial frameworks and are discovering through incidents that “partial” means the ungoverned population is larger than they realized.

The formal process picture completes this view. 76.3% currently have a process while 19.8% say one is still in development. Only 29.2% say it is always followed, while 47% acknowledge inconsistent adherence. A governance process followed inconsistently is not a governance control – it is the organizational record of an intention that agents are routinely entering production without satisfying.


How integrated is AI agent governance with your broader IT and security governance processes (e.g., change management, access reviews, risk assessments)?



Does your organization have a formal, documented process for moving agentic AI solutions from prototype to production?




You indicated that your organization has a formal process for moving agentic AI from prototype to production. How would you rate your organization’s adherence to that process?




The downstream consequence is visible in how security reviews interact with deployment timelines. Nearly 68% of respondents have had a security, identity, or governance review directly delay, scale back, or stop a project – 32.2% more than once. This is not primarily evidence that controls are working. It is evidence that controls are arriving after deployment decisions have already been made, generating friction rather than risk reduction. The organizations in which security reviews accelerate deployment rather than block it are those that embed governance from the design stage – not those that position it as a final approval gate.

This late-gate dynamic has a compounding organizational effect. Security reviews that arrive after a project has been built, staffed, and approved by business leadership create pressure – not to fix the governance gap, but to route around the review or change the security/governance requirements. Over time, teams learn to avoid triggering formal security checkpoints through informal deployments, shadow agent inventories, and production deployments that live in the governance record as “still in pilot” long after they are running at meaningful scale. The gap between what the formal process requires and what actually happens is not a governance intent problem. It is a governance architecture problem – and the data shows that 47% of the sample is living with its consequences.

Has a security, identity, or governance review ever been the direct reason an agentic AI project was delayed, scaled back, or stopped?






Agents are Already Acting Outside Their
Intended Scope

Enterprise AI incidents have been largely invisible in the public record – undisclosed, uninvestigated by external parties, and absent from the regulatory enforcement and press coverage that would make them visible. The survey data suggests that invisibility is not because these incidents are rare. It is because they are substantially underreported. Among the enterprises actively deploying agentic AI, the experience of an agent acting outside its intended scope is not the exception. It is the majority experience.

When EMA asked whether an AI agent had ever taken an action outside its intended scope or authorized access, 29.2% reported a confirmed incident with measurable organizational impact – data exposure, financial consequences, operational disruption, or reputational damage. Over 35% reported a caught near-miss. Only 34.2% reported no known out-of-scope agent action.


The combined rate (nearly 65%) substantially exceeds what the public discourse around enterprise AI risk suggests. For a technology that most organizations in this sample have been running in production for less than two years, that baseline should reframe how the industry talks about AI risk. The question is no longer whether enterprises will experience an AI agent incident. For the majority of active deployers in this sample, that question has already been answered. The questions that remain are what prevents incidents from occurring, how quickly incidents are detected, how effectively they are contained, and whether the evidence exists to reconstruct what the agent actually did.



The detection picture is as consequential as the incident rate itself. Only 49.5% of respondents identify automated monitoring as their primary mechanism for detecting out-of-scope agent behavior. Thirty-two percent only discover rogue agent behavior during an internal review or audit – meaning the agent continued operating outside its intended scope until a scheduled checkpoint. Three and a half percent – seven organizations in a sample of 202 – first learn about out-of-scope agent behavior when a customer or external partner reports it.

That 3.5% figure is the starkest data point in the dataset. The incident has progressed to external visibility before anyone inside the organization has detected it through any internal process. The security team learns about a rogue AI agent the same way it might learn about a data breach from a journalist – from the outside, after the fact, with no opportunity to limit the damage before it becomes visible to the people on whose trust the organization depends.


When an AI-driven or agentic action behaves outside its intended scope (including incorrect or harmful outcomes), how does your organization most often first become aware of it?




Detection is one half of the incident response equation. Containment is the other – and the data is equally concerning. Only 32.2% can detect and contain an out-of-scope agent action within minutes through automated mechanisms. Almost 55% could detect within hours, but required manual steps to contain it. Eight percent would only identify the issue during a scheduled review. Four percent would only know after measurable impact had already occurred. Given that AI agents operate at machine speed, anything slower than real-time detection of out-of-scope agent actions is a serious problem.

A second question asked specifically how quickly organizations could intervene to restrict or suspend an agent’s unexpected activity. Forty-eight percent would require hours of manual effort. Nearly 10% would need days. An agent operating in an IT help desk, security operations, or customer-facing workflow can take hundreds of actions in the hours between anomalous behavior beginning and a human reviewer initiating containment. Manual containment is a defensible posture for a handful of closely monitored pilots. It is not defensible for the multi-agent production portfolios most enterprises in this sample are currently operating.

If an AI agent began acting outside its intended scope in production right now, how quickly could your organization detect the issue and contain it?




If an AI agent begins behaving unexpectedly or outside its intended scope, how quickly can your organization intervene to restrict or suspend its activities?




Once an incident is contained, the organization faces a second challenge: understanding what the agent actually did. Forty-six percent of respondents cannot easily and completely produce an audit trail of all actions taken by a specific agent over the past 30 days. For nearly half the sample, the reconstruction required after an incident – what the agent accessed, what data it touched, what actions need to be reversed or remediated – requires significant manual investigation, if it is possible at all. The incident record is real. The containment posture is insufficient. For nearly half of active deployers, the audit infrastructure needed to understand the scope of what happened does not exist.

A separate but related risk concerns what happens to agents that are no longer actively managed. Respondents reported that agentic AI pilots in their organizations have been paused indefinitely (18.8%), formally discontinued (11.9%), or restarted from scratch after abandonment (12.4%). These were real deployments – provisioned with credentials, granted access to production systems, and integrated into operational workflows. When those pilots were discontinued, many were not cleaned up.

If your organization needed to, could you produce an audit trail of all actions taken by a specific AI agent over the past 30 days?




Of the pilots that did NOT reach production, what happened to most of them?




The leading stall factors tell their own story. Security risk concerns were cited as a major factor in 48.5% of stalled pilots and identity and access management gaps in 22.3%. The agents most likely to have been discontinued because their governance posture was inadequate are precisely the agents most likely to have been decommissioned without the credential revocation, permission cleanup, and integration disabling that clean decommissioning requires.


The accountability picture for these incidents is equally unresolved. When an agent produces a harmful outcome, survey data shows responsibility is distributed almost evenly across four functions: the deploying team, a dedicated AI team, central IT, and GRC, with only 2% acknowledging no clearly defined owner. The result: 98% believe someone is accountable, but the distribution is diffuse enough that rapid, coordinated response is harder than it needs to be. Diffuse accountability and slow manual containment are a combination that allows incidents to expand before a coherent response begins.

Thinking about agentic AI pilots that stalled, were paused, or were abandoned, which of the following were major factors?



When an AI-driven or agentic action produces an incorrect or harmful outcome, where does primary accountability sit in your organization?






The Governance Deficit: Why Agents Act Out of Bounds, and What Explains It

The 65% incident and near-miss rate is not random. It is the product of a compound governance failure – a series of gaps in how agents are provisioned, how their access is reviewed at runtime, how their behavior is monitored, and who in the organization is actually accountable for what they do. No single gap explains the incident rate. All of them, operating together in most enterprises in this sample, produce it.


When respondents were asked how confident they are that their AI agents do not have more access than they need, 48.5% said very confident and 45.5% said somewhat confident. Ninety-four percent of the sample expressed at least moderate confidence that agents are not overprovisioned.

That confidence exists in the same sample where only 32.7% of respondents actually provision agents with least-privilege access. The remaining 67.3% are using some form of broader, less controlled, or unreviewed access: 38.1% use broad-standing access reviewed periodically, 7.9% use broad access rarely or never reviewed, 10.9% provision agents with the full permissions of the deploying user, and 8.9% have no consistent process at all.

How confident are you that AI agents in your organization do NOT have more access than they need to perform their intended tasks?



When a new AI agent is provisioned, what access does it typically receive?




A similar and equally troubling confidence gap recurs at the governance framework level. Ninety-two percent express confidence that their existing governance and compliance frameworks can keep pace with the rate of change in agentic AI. This is the same population in which 67.8% have had a security review directly delay or stop a project – evidence not of frameworks keeping pace, but of governance arriving late and creating friction. High confidence in a framework that is structurally positioned as a late-stage gatekeeper is not a sign of mature governance. It is a sign that the framework’s limitations have not yet been tested at the scale that is coming.

Organizations that believe their governance posture is adequate do not look for evidence that it is not. This confidence gap is consequential beyond its own measure. It represents a posture that has three predictable consequences:

- They do not look for evidence that something is wrong until an incident, an audit, or a customer report makes it undeniable
- They do not invest in the automated monitoring that would surface behavioral anomalies because the assumption is that agents are operating within their intended scope
- They do not conduct the access reviews that would reveal overprovisioning because the reviews feel unnecessary when confidence is high


The confidence gap is not a calibration error. It is the mechanism through which the other gaps go unaddressed. It serves as the most direct explanation for why agents take unauthorized actions.



The most direct explanation for why agents take unauthorized actions is that authorization is evaluated at the wrong time (or not at all). Only 34.2% of respondents evaluate agent authorization at the moment an agent attempts to take a specific action. The majority rely on periodic policy reviews or standing permissions – access that was scoped at some prior point and has remained in place since.


The consequence is compounding. An agent provisioned with access to a customer database for one task retains that access across every subsequent task, whether or not those tasks require it. Standing permissions, never revalidated at the moment of action, become permanent entitlements. Over time, the gap between what an agent is technically authorized to do by its standing permissions and what it actually needs for any given task grows wider, and that widening gap is the direct mechanism through which overprovisioning becomes the 65% incident rate.

When an AI agent attempts to perform a business or operational action, how is authorization typically determined?



Effective governance of agent actions requires that authorization be scoped to the specific task, evaluated at the moment of execution, and defined within explicit behavioral boundaries – not granted broadly at provisioning time and left to accumulate. The organizations building that capability are the ones with agents that are operating with defined job descriptions rather than inherited access portfolios. Identity is one of the contributing causes of the agent action problem, not a separate issue. Fifty-four percent require and enforce unique, individually identifiable identities for all AI agents, but 32.2% require it without consistently enforcing the mandate in practice. An additional 8.9% are developing the requirement, and 3.0% acknowledge that agents share or inherit credentials from user or service accounts.


The 32.2% “required, but not enforced” population is the clearest expression of the policy-practice gap: organizations that have written the correct governance requirement and have not completed the operational work of making it real. Without enforced unique identity, the other controls – access reviews, audit trails, behavioral monitoring – are built on a foundation that does not reliably know which agent is which. When an incident occurs, attributing the action to a specific agent, with a specific authorization state, is significantly harder when that agent shares credentials with others or with a human account.



The governance gaps described do not exist in a vacuum. They exist in a structural context in which security expertise is systematically absent from the decisions that create the most risk. Fifty-four percent of respondents named the CIO or CTO as the production approval authority for agentic AI deployments. The CISO or CSO holds that role for only 15.8%. Head of AI or AI CoE accounts for 11.9%.

When security leadership is a secondary voice in deployment decisions, security review arrives after the decision, producing friction and delay rather than risk reduction. The organizations in which security and governance reviews accelerate deployment are those that bring security into the design conversation, not the approval gate. Provisioning, behavioral boundaries, monitoring architecture, and incident response planning belong at the pilot stage, not the production approval stage.


Who in your organization approves agentic AI deployments into your production environment?




The Model Context Protocol (MCP) has become the de facto mechanism for connecting AI agents to external tools, data sources, and APIs. Almost 14% of respondents allow MCP connections without any restriction. Fifty-six percent limit connections to an approved list, but of those, only 49.1% have a dedicated governance team maintaining that list on a regular cadence. Just under 40% rely on ad hoc IT review. Three percent rely on individual team self-reporting with no central oversight.

The runtime authorization gap connects directly to MCP: connections approved for one task become standing entitlements, never revalidated as agent tasks evolve. An agent permitted to connect to an external data service for one use case carries that connection into every subsequent use case whether or not it is appropriate, and whether or not anyone has reviewed whether the original approval still reflects current practice. MCP governance is not a separate problem. It is the external connectivity expression of the same compound failure that runs through provisioning, identity enforcement, and runtime authorization.

Does your organization currently permit employees to connect AI agents to external tools or data sources via protocols such as Model Context Protocol (MCP)?



You indicated that your organization limits AI agent connections to external tools to an approved/vetted list. Who is responsible for maintaining and auditing that list?





Conclusion

The data in this research tells a story about control, or the lack of it.

Enterprises are already running AI agents in production and core business functions, and already doing things that they were never authorized to do. The majority of organizations in this survey that have deployed AI agents have experienced a lack of control. Most took hours to respond. Many still cannot fully explain what happened. None of these enterprises can claim they were not aware. Every organization knows that agentic AI carries risk. Most have policies and most have expressed confidence in their policies and processes.

The failure is operational – the gap between what organizations have written down and what is actually enforced, monitored, and actionable when an agent goes out of bounds.

When viewed in total, the data tells an obvious story:

- Deployment and use have outrun the infrastructure built to govern them
- Agents are acting outside their intended scope at a rate the industry has not yet acknowledged
- Most organizations cannot detect it quickly, stop it fast, or reconstruct it afterward

The combined failures that explain these are not isolated or independent problems. They are the same problem at different points in the agent lifecycle. The practitioners who have already been through an incident are nearly unanimous in their guidance: build the controls before scaling, not after the first event forces the conversation.

Three actions the data supports most directly:

Govern what agents do, not just what they are.

The most direct cause of the incident rate documented in this research is that agents routinely act on access never scoped for their current task, with no authorization check at the moment of action. Defining what an agent is permitted to do – task by task, within enforced behavioral limits – is the single most consequential control available. Treating agent authorization as a provisioning event rather than a runtime discipline will continue to produce out-of-scope actions discovered only after the fact.

Build detection and containment capabilities before expanding agent deployment.

Most enterprises in this sample rely on manual processes, measured in hours, to detect and respond to out-of-scope agent behavior – and some learn about it only from customers. That posture cannot scale alongside multi-agent production portfolios. Automated behavioral monitoring, defined response playbooks, and rapid containment mechanisms belong in the deployment prerequisites, not in the post-incident improvement plan.


Treat agent decommissioning as a security discipline, not a project closure checkbox.

Nearly one-third of agentic AI pilots in this sample have been paused, discontinued, or abandoned – many without the credential revocation and permission cleanup that clean decommissioning requires. Every discontinued pilot retaining active credentials is an ungoverned exposure. As discontinuation rates rise with the first wave of AI investment renewals, organizations without enforced decommissioning standards will find their residual attack surface growing faster than their active deployment footprint.




Research Methodologies and Demographics

Which of the following best describes the department or functional area in which you work?




You indicated that your department is IT-related. Which of the following BEST describes your specific role?




Sample Size = 202

In total, how many employees are currently working in your organization? Select one.




Sample Size = 202

In which region is your organization’s headquarters located?



Which of the following best describes your organization’s primary industry?



Sample Size = 202






Christopher M. Steffen, CISSP, CISA
Vice President of Research
CSteffen@enterprisemanagement.com



Check out the CyberSecurity Awesomeness Podcast:

<https://www.cybersecurityawesomeness.com>

CYBERSECURITY
AWESOMENESS PODCAST





About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Our team of analysts combines practical experience with a deep understanding of industry best practices and emerging vendor solutions to help clients achieve their strategic objectives. Learn more about EMA research, analysis, and consulting services at www.enterprisemanagement.com or follow EMA on [X](#) or [LinkedIn](#).

This report, in whole or in part, may not be duplicated, reproduced, stored in a retrieval system or retransmitted without prior written permission of Enterprise Management Associates, Inc. All opinions and estimates herein constitute our judgement as of this date and are subject to change without notice. Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies. "EMA" and "Enterprise Management Associates" are trademarks of Enterprise Management Associates, Inc. in the United States and other countries.

©2026 Enterprise Management Associates, Inc. All Rights Reserved. EMA™, ENTERPRISE MANAGEMENT ASSOCIATES®, and the mobius symbol are registered trademarks or common law trademarks of Enterprise Management Associates, Inc.