

Prepared for



AI On Your Terms

October 2026 EMA White Paper

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

Table of Contents

- 1** Introduction
- 2** Defining AI on Your Terms
- 2** Defining an Agentic SOC
- 3** Deployment Control: Where AI Runs
- 4** Data Control: Sovereignty as a Prerequisite
- 6** Autonomy Control: Visibility Over AI Decision-Making
- 8** The Payoff: What AI on Your Terms Actually Buys You

Introduction

The evolution has been frantic: AI-driven threats have exponentially increased through agentic systems and everyone – attackers and defenders alike – are automating at machine speed. The rate that organizations adopt AI is restrained by human-in-the-loop processes instead of technical evaluations of actual capabilities. Almost 82% of organizations report diminished AI deployment plans directly tied to unresolved friction, and 66.7% see delays of a month or longer before an AI security initiative ever reaches production.¹ This is a continued failure of human processes/controls, not technical capabilities, and decision makers often confuse the two for the inability to realize productivity gains.

How long were your AI projects delayed from moving into production due to data security or compliance reviews?



AI's payoff in the security operations center (SOC) is real, but only realized when security teams have the ability to control where AI runs, what data AI touches, and how much decision-making authority AI is allowed to have.

We call that combination AI on your terms.

¹ EMA, "State of AI Friction" 2026

Defining AI on Your Terms

AI on your terms is not a punchy marketing phrase. It is an organizational philosophy made up of three components that define what responsible AI adoption actually requires inside security operations.

Starting with AI: Traditional machine learning and automation scores and classifies. Generative AI algorithmically generates content and conversation as requested, in an easily consumed matter.

Agentic AI is different: it examines a request, generates an outline or plan for solution, takes multiple steps to address, and (when properly configured and permissioned) can reach into the various tools and systems of the infrastructure with limited human intervention and a certain level of autonomy to continue the request.²

There are three components that make AI on your terms from there:

- Where the AI runs...on-premises, air-gapped, or in the cloud, and who actually decides that.
- What data it touches, as well as what an agent sees and processes while doing its job.
- How much autonomy it holds...the human-in-the-middle debate, what AI is expected to decide on its own, and how much visibility a human retains into why AI did what it did.

For clarity: none of this is an argument against AI or augmenting a SOC with agentic AI usage, but only against the AI black box version of it. Security and GRC teams have spent the last several years hearing “trust us” from AI vendors, only to see that trust violated continuously and constantly (to the point where many business executives are purposely sourcing solutions without AI components due to potential risks and vulnerabilities that AI brings to their environments).

Defining an Agentic SOC

AI on your terms is the philosophy while the agentic SOC is where it's applied.

An agentic SOC is a security operations model in which purpose-built AI agents carry out the core work of investigation, remediation, and threat hunting, operating within human-defined guardrails and a level of authority that the organization expands as trust in the system is earned. It is a part of the organization's security program, not a product that one can buy from a vendor booth at a conference, with agentic AI as one part and the strategy, processes, governance, and oversight around them as the rest.

² EMA, “What Defines an AI or Autonomous SOC?” 2026

The agentic SOC is different than what came before it in three significant ways:

- Agentic AI for the SOC has the ability to reason instead of just following scripts or playbooks. Where a SIEM/SOAR generate alerts or runs pre-scripted playbooks, the agentic AI work out what context they need, query for it, and connect low-level signals into an evidence-backed verdict.
- The Agentic SOC MUST be transparent by design. Every action/recommendation/verdict has to be explainable and auditable to analysts, leadership, and regulators.
- Authority is tiered and revocable (and – ideally – provisioned “just-in-time”). Actions run as fully autonomous, human-in-the-loop, or manual only, scoped by role and gated by confidence, with a way to halt everything at once.

Deployment Control: Where AI Runs

Cloud-only AI security tools make up a significant share of the market and the reasons are obvious: frontier models run in the cloud, and they make it easy for security vendors to offload logic into AI with prompts (NOTE: regulatory requirements for AI data use are rapidly evolving, even as this paper is being drafted, and changes are being suggested and legislated).

Research from EMA shows that 47.4% of security leaders reported approving AI security initiatives at their organization, and that approval rate tracks closely with architectural and deployment concerns, not skepticism about whether AI works.³ The AI is usually not the problem, but rather where the AI is going to run.

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



³ EMA, DSPM (Data Security Posture Management) research, 2026

Most GenAI (generative AI) security tooling was created with a cloud-first mindset, with on-premises or air-gapped deployment treated as an afterthought (if it was considered at all). That choice is exactly the gap that stalls adoption once a security team tries to move an initiative past architecture review. A regulated bank, a defense contractor, or a hospital system does not get to waive its data residency obligations because a vendor's roadmap assumed public cloud by default.

Some solutions already support on-premises, air-gapped, and cloud deployment for this exact reason. The move toward a dedicated private cloud extends that principle into organizations in which public cloud deployment would never have cleared an architectural review.

It is the same problem that has been front and center in the cybersecurity industry for years: vendors are “great” at the detection aspect of the problem, but not as “great” at the response or performance of their solution to actually solve – in this case, on whether their AI can run entirely inside a customer's own walls.

Data Control: Sovereignty as a Prerequisite

If the rapid adoption of AI has taught us anything, it is that we are woefully unprepared for the increases of data that AI solutions produce, what happens to it, and what IT and security leaders are supposed to do with it all. Sovereignty has to mean more than “where is the data stored?” EMA research recently found that roughly 80% of organizations require or strongly prefer in-region scanning of their data, while only 37.3% consider their data governance maturity adequate to the task.⁴


Storage location is the easy half of the sovereignty question. The harder part is controlling what an AI agent actually sees and processes once it is operating inside the environment. EMA research identified eleven distinct categories of data that agentic systems routinely touch, from operational telemetry to credentials to business logic, and built a four-pillar governance framework around a simple observation: most organizations have never fully inventoried what their AI agents can reach.⁵

⁴ EMA, “The Invisible Data Problem” (agentic AI data catalog white paper), 2026

⁵ EMA, Agentic AI and identity security survey, 2026

“Where is my data?” and “what can my AI agent see?” may sound like the same question, but experience (mixed with a healthy dose of reality) has proven that they are not. When treated identically, organizations may end up compliant on paper, but non-compliant in practice.

Even among organizations that have made real data protection investments, 64.4% use a commercial DLP (data loss prevention) solution and 57.4% enforcing data classification and tagging at runtime, a full 11.9% still run AI agents with no exfiltration-specific controls at all.⁶



While most organizations are trying to get a handle on their agentic data, most have not finished, and the work will evolve almost as quickly as the agentic solutions themselves. It is an ongoing architecture decision and that is how AI on your terms treats it.

⁶ EMA, Agentic AI and identity security survey, 2026

Autonomy Control: Visibility Over AI Decision-Making

AI on your terms is not only a philosophy, but a preference. When organizations were asked if they wanted a “trust but verify” approach, human oversight and approval required initially with autonomy expanded over time, or immediate autonomous operation with minimal human involvement, 74.6% of organizations chose trust but verify. Only 17.2% wanted immediate, full autonomy.⁷

That preference shows up twice, independently, in two separate EMA studies surveying two different populations. Research into agentic AI and identity found adoption is already real and moving fast: 87.9% allow AI agents today in some form, though nearly half of that figure is only pilot-stage, and 94.9% expect gen AI and agentic AI to meaningfully increase their level of automation.⁸ The appetite is not in question. The caution is deliberate, and two separate respondent pools landing on the same instinct is a stronger signal than either study alone.

The caution is not hypothetical, nor without valid reasons. That same research found 29.2% of organizations report an AI agent has already taken an action outside its intended scope with measurable impact, financial, operational, reputational, or data exposure, and another 35.6% caught a near miss. Roughly two out of three organizations already have a live incident or a close call on the books.⁹


⁷ EMA, “The Mainframe in 2026” survey, 2026

⁸ EMA, Agentic AI and identity security survey, 2026

⁹ Ibid

Part of the explanation is how agents get provisioned in the first place. Just 32.7% of organizations give new AI agents strict least-privilege access. The rest hand agents broad standing access: reviewed periodically at best (38.1%), rarely at all (7.9%), or inherited wholesale from the deploying user's own permissions (10.9%). Call it the keys-to-the-kingdom problem wearing an AI badge instead of a human one.¹⁰

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



That gap carries a business cost, not only a security one: 67.8% of organizations say a security, identity, or governance review has directly delayed, scaled back, or stopped an agentic AI project at least once, 32.2% more than once.¹¹

Ask those same organizations what would most accelerate getting agentic AI into production, and the top answers are not “a better model.” They are better monitoring and audit capabilities, better visibility into what agents access and do, and stronger identity and access controls.¹² In customers' own words, that is a fairly precise description of what AI on your terms delivers.

¹⁰ EMA, Agentic AI and identity security survey, 2026

¹¹ Ibid

¹² Ibid

The Payoff: What AI on Your Terms Actually Buys You

Control is not the opposite of speed. It is what makes speed sustainable, and two independent EMA studies arrive at the same conclusion from different directions. EMA research found that unresolved friction drags down AI deployment plans for 81.6% of organizations and adds month-plus delays for 66.7% of them.¹³ A separate EMA study also found that a security, identity, or governance review has directly delayed, scaled back, or stopped a project for 67.8% of organizations. Different surveys but the same conclusion: the absence of control is not a safety trade-off against speed.¹⁴ It is what is actually slowing everyone down.

While the opposite is not always better, in this case, the payoff is straightforward. Organizations that can answer “where does it run, what does it touch, how much authority does it have” with architecture and process can shorten the approval cycle, sidestepping the stalled/abandoned pilot pattern that roughly one-third of agentic AI projects fall into (and potentially even help eliminate another significant problem not directly discussed: the use of shadow AI).

In the SOC, those same types of controls are the only realistic path to combat alert-fatigue that agentic SOC has promised since the beginning. Analysts should not and will not trust an opaque agent’s triage decisions, but can grow to trust a transparent one.

AI on your terms is not a slower route to that outcome. It is the only route that actually arrives.

¹³ EMA, “State of AI Friction” 2026

¹⁴ EMA, Agentic AI and identity security survey, 2026



30
YEARS

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.