ISSUE 001, JUNE 2026

THE AI RUNTIME REPORT

Everyone Wants AI Agents. Few Are Asking What Happens Next.



Every AI connection expands capability.
Every AI action expands responsibility.

PUBLISHED BY EVE SECURITY

Everyone Wants AI Agents.
Few Are Asking What Happens Next.

Artificial intelligence has entered a new phase. Not long ago, AI was used to answer questions, summarize documents, and generate content. Today, it's becoming an active participant in the enterprise.

Employees are connecting AI to email, CRM, HR systems, cloud platforms, and internal knowledge bases to complete work that once required a person.

The benefits are significant. Faster execution. Better decisions. Less repetitive work.

But every new connection gives AI more context, more access, and more responsibility. That raises an important question.

What happens after an AI agent is deployed?

What organizations gain

Faster execution

AI completes routine work in seconds.

Better context

Information from multiple systems comes together automatically.

Smarter decisions

AI identifies patterns people might miss.

Greater scale

One assistant can support thousands of users.

AI no longer just answers questions.
It takes action on your behalf.

The question isn't whether AI can help your business. It's whether you understand what it's doing once it can act.

Every AI Connection Expands Responsibility.

AI agent

Cloud	CRM	Email
AWS, Azure, Google Cloud	Salesforce, Dynamics	Gmail, Outlook
Connected	Connected	Connected

HR	Development	Finance
Workday, BambooHR	GitHub	SAP, NetSuite
Data access	Connected	Restricted

Current activity

Actions

- Reviewing customer history
- Reading recent emails
- Checking support tickets
- Drafting executive briefing
- Preparing meeting notes

Status

Tasks in progress	5
Systems connected	6
Policies applied	12
Actions today	23

> Every AI connection expands capability. Every AI action expands responsibility.

No One Manually Opened Those Applications. The AI Did.

At 9:00 a.m., an employee asks an AI assistant to prepare for a customer meeting. Within seconds, the assistant searches Salesforce, reviews recent emails, summarizes support tickets, checks open opportunities, and drafts an executive briefing.

01

9:00 AM

Employee asks AI assistant to prepare for a customer meeting.

02

9:01 AM

AI searches Salesforce for account activity and open opportunities.

03

9:02 AM

AI reviews recent emails in Outlook.

04

9:03 AM

AI reviews open support tickets and recent customer interactions.

05

9:04 AM

AI drafts executive briefing and meeting agenda.

Connected systems

Salesforce, Outlook, Support Platform, Docs & Drive, Calendar.

Objective

Prepare for customer meeting. Completion: 100%.

> **No human touched those systems. The AI did.**

This is no longer experimental technology. Across nearly every industry, AI is moving beyond answering questions and beginning to perform work.

Every new connection makes AI more capable. It also increases the need for visibility into what those systems can access and the actions they take.

AI No Longer Just Answers Questions. It Takes Action.

01

AI understands context

AI no longer works from a single prompt. It reasons across conversations, documents, and business data to understand intent and deliver relevant results.

02

AI connects to enterprise systems

Email, CRM, HR, cloud platforms, code repositories, collaboration tools, and databases are becoming part of the AI workflow. The connections are growing every day.

03

AI takes action

Modern AI doesn't just recommend what to do. It can complete tasks, initiate workflows, update records, and make decisions within defined boundaries.

04

AI requires runtime visibility

Once AI is connected and acting, traditional security stops at the door. Organizations need real-time visibility into every action to manage risk and maintain trust.

> AI no longer just answers questions. It takes action.

The AI Race Is Accelerating. Security Must Keep Pace.

Over the past year, nearly every major technology company has shifted from building AI assistants to building AI that can reason, connect to enterprise systems, and take action.

The announcements are different, but the direction is the same: AI is becoming an active participant in the enterprise.

Recent developments. Bigger implications.

Google

Gemini + Workspace

AI can read email, summarize meetings, draft documents, access Drive, and automate everyday work.

SECURITY QUESTION

What enterprise data can those agents access?

Microsoft

Copilot

AI is becoming part of Word, Excel, Outlook, Teams, Dynamics, GitHub, and Windows.

SECURITY QUESTION

What actions should AI be allowed to perform without human approval?

OpenAI

Operator & Agents

The focus is shifting from conversations to autonomous task execution.

SECURITY QUESTION

How do you govern actions instead of prompts?

Anthropic

Model Context Protocol (MCP)

Standardizing how AI agents connect to external tools and business systems.

SECURITY QUESTION

Can every connection be trusted?

AWS

Amazon Bedrock AgentCore

Helping organizations build enterprise AI agents with memory, tools, orchestration, and governance.

SECURITY QUESTION

Who decides what an agent is allowed to do?

Salesforce

Agentforce

CRM agents capable of interacting with customers and executing business processes.

SECURITY QUESTION

How do you verify every customer-facing action?

> The question is no longer whether AI can take action. The question is whether you can see, understand, and govern those actions in real time.

One Pattern
Connects Them All.

01 **Smarter AI**
Models reason better and understand more context.

02 **More connected AI**
AI connects to more tools, systems, and data sources.

03 **More autonomous AI**
AI moves from assisting to acting on behalf of users.

04 **More business value**
Faster execution, higher productivity, new capabilities at scale.

05 **More risk**
More access, more complexity, more potential for harm.

06 **Greater need for runtime security**
Visibility, context, and control at the moment of action.

As AI becomes more autonomous, runtime visibility and control are no longer optional. They are essential.

Watch list

MCP adoption
As more vendors support MCP, organizations will need stronger controls around tool access and trust boundaries.

AI identity
Expect every AI agent to have its own verifiable identity, just like users and service accounts today.

Agent-to-agent communication
Autonomous systems will increasingly coordinate with one another, creating new governance and visibility challenges.

AI regulation
Governments and industry bodies are beginning to define expectations for transparency, accountability, and auditability.

AI Runtime Security Secures the Moment That Matters Most.

AI agents can reason, access data, connect to tools, and take action across your systems.

AI Runtime Security focuses on the only part of an AI interaction that can change your business: **the action.**

It evaluates intent, context, risk, and policy in real time and decides what should happen next.

AI Runtime Security is the decision layer for autonomous AI. It secures every action at the moment that matters most.

The four decision points

01 What is the AI trying to do?

Understand intent, goal, and the target.

02 Is this action appropriate?

Evaluate risk, policy, and business context.

03 Should we allow, block, or require approval?

Make the right decision in real time.

04 Record and explain.

Log the action, outcome, and rationale for audit and learning.

Why it matters

Real-time visibility

See every AI action as it happens across systems, users, and tools.

Context-aware control

Decisions adapt to user, data, intent, environment, and business rules.

Risk reduction

Prevent inappropriate actions before they cause impact.

Auditability & trust

Every action is logged, explained, and traceable for complete accountability.

Understand. Evaluate. Decide. Enforce & Record.

01 **Understand**

The AI agent makes a request. We identify the intent, target, and context.

02 **Evaluate**

We analyze the action using policy, risk signals, user behavior, data classification, and environment.

03 **Decide**

We determine whether to allow, block, or require approval, based on real-time policy and business rules.

04 **Enforce & record**

We enforce the decision instantly and log the outcome, rationale, and full context.

05 **Secure outcomes**

The right action happens at the right time, driving innovation with confidence.

Runtime decisions in action

AI ACTION & CONTEXT	RISK	DECISION	WHY IT MATTERS
Summarize customer email thread Sales rep preparing for renewal call	Medium Contains PII	Allow Within policy	Helps rep respond faster and accurately
Download Q4 financial report Analyst requesting restricted file	High Confidential data	Require approval Manager validation	Ensures sensitive data is reviewed first
Update customer record Support agent resolving ticket	Low Standard data	Allow Within role	Keeps information accurate
Send follow-up email Account exec outreach	Medium External recipient	Allow Approved template	Consistent communication at scale
Delete inactive user account IT admin offboarding user	High Irreversible action	Block Policy violation	Prevents accidental loss of access

> Traditional security decides who gets access. AI Runtime Security decides what happens next.

Before You Deploy an AI Agent, Ask These Five Questions.

Every AI agent introduces a new decision maker into your environment. Before moving from pilot to production, every security leader should be able to answer five questions with confidence.

01

Who owns this agent?

Every AI agent should have a clearly defined owner responsible for its purpose, behavior, and business outcomes.

- Who is accountable for the agent's purpose and performance?
- Who approves changes to its connections?
- How is ownership maintained over time?

No owner. No production deployment.

02

What can it access?

Every connection increases both capability and risk. Look beyond permissions and understand the business impact.

- What applications can it reach?
- What sensitive data can it access?
- What actions can it perform?
- What systems should always remain off limits?

Every connection expands opportunity and exposure.

03

How does it make decisions?

Identity tells you who initiated the request. Context explains why the action is happening. Modern AI decisions should never be a black box.

- What objective is the AI trying to accomplish?
- What information influenced the decision?
- Which policies were evaluated?
- Can those decisions be explained later?

Intent matters as much as identity.

> Strong security decisions today lead to strong AI outcomes tomorrow.

Visibility and Control. Before the Action Completes.

04

Can you see every action?

Security visibility should extend beyond prompts and into runtime. Every action should be observable.

- What action occurred?
- Which systems were accessed?
- Which user initiated it?
- What data was involved?
- Was the action successful?

If you can't see it, you can't govern it.

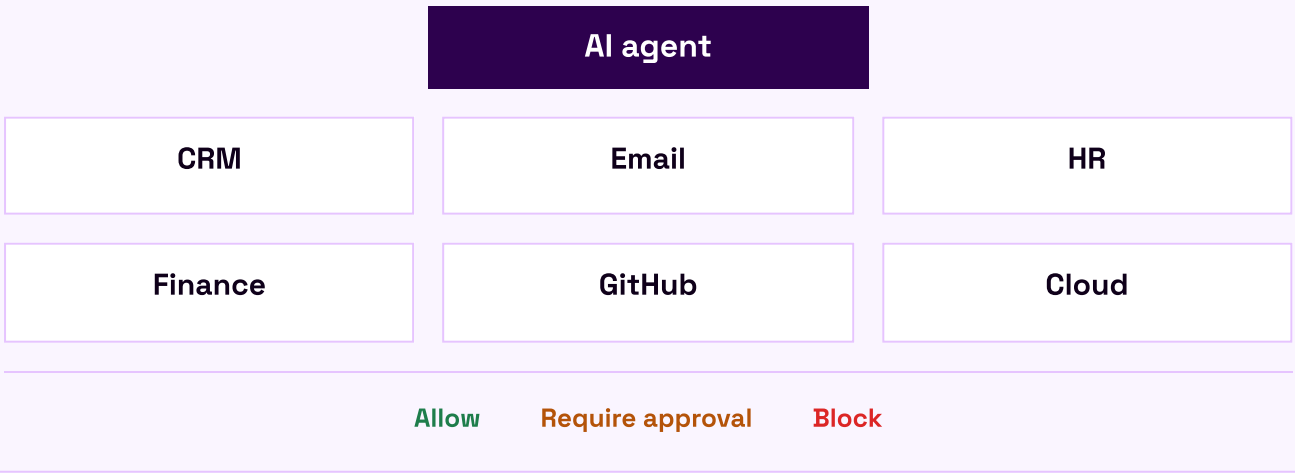
05

Can you intervene in real time?

The most valuable security decision is the one made before damage occurs. Runtime controls should evaluate every action before it executes.

- Can actions be blocked?
- Can approvals be required?
- Can policies adapt to changing context?
- Can risky behavior be stopped immediately?

Control only matters before the action completes.



> Control only matters before the action completes.

AI Agent Production Readiness Score.

AI agents can drive enormous value, but only when governed with clarity, guardrails, and runtime control. These five questions help reduce risk, build trust, and prepare your organization for what comes next.

Score yourself. Answer yes or no for each statement before moving an agent into production.

SCORE YOURSELF	YES	NO
Every agent has an owner	☐	☐
Every connection is understood	☐	☐
Every decision is explainable	☐	☐
Every action is visible	☐	☐
Every action can be controlled	☐	☐

5/5

Ready for production. You have the foundation for secure AI at scale.

3 to 4/5

Ready for limited deployment. Close the remaining gaps before scaling.

0 to 2/5

You're deploying AI faster than you're governing it. Significant risk to your business.

Visibility across systems and actions is the foundation of AI risk management.

The organizations that succeed with AI won't be the ones deploying the most agents. They'll be the ones asking the right questions before those agents take action.

> Production readiness: If you can't confidently answer these five questions today, there's still work to do before autonomous AI belongs in production.

The Future Belongs to Organizations That Can Trust What AI Does.

> AI will continue to evolve, smarter, more connected, and more autonomous.
>
> The gap between innovation and governance will only widen.
>
> The organizations that thrive won't be the ones that move the fastest.
>
> **They'll be the ones that see clearly, make better decisions, and act with confidence.**

The path forward

01 **See clearly**
Gain real-time visibility into what AI agents can access and what they are doing.

02 **Set guardrails**
Define policies based on risk, intent, and business context.

03 **Decide in real time**
Evaluate every action as it happens and respond instantly.

04 **Learn & improve**
Capture every decision to strengthen models, policies, and outcomes.

05 **Build trust**
When AI acts within guardrails, the entire organization can move forward with confidence.

The organizations that win with AI won't be the ones that deploy the fastest. They'll be the ones that trust what their AI is doing.

About

The AI Runtime Report tracks how production agents are governed, secured and kept accountable at the point of decision, one issue at a time.

This is Issue 001, June 2026.

We're building a future where AI can be powerful and safe.
AI Runtime Security is how we get there.

Thank you for reading.

Follow Eve Security on LinkedIn

PUBLISHED BY EVE SECURITY

Gain Visibility.
Restore Trust.
Enforce Control.
Accelerate AI Securely.