ISSUE 002 · JULY 2026

THE AI RUNTIME REPORT

You Can't Secure What You Can't See.

Why AI visibility is becoming every CISO's biggest blind spot.

AI ACTIVITY VISIBLE **12%**

Most of what your agents do runs unseen.

PUBLISHED BY EVE SECURITY

01

FEATURE STORY

The AI Visibility Gap

Every major evolution in cybersecurity has started with visibility. AI Runtime Security is no different.

Users | Devices | Servers | Applications | Cloud Workloads | **AI Agents**

Firewalls gave us network visibility. EDR gave us endpoint visibility. CNAPP gave us cloud visibility. **Each made policy, monitoring, and response possible.**

Now employees create assistants, developers deploy agents, and business units connect AI to enterprise systems in days. Most security teams still have no complete inventory of what exists.

FIVE QUESTIONS WITHOUT ANSWERS

01 How many AI agents exist?

02 Who created and owns them?

03 What systems can they access?

04 What data are they using?

05 What actions are they taking right now?

FROM SIGHT TO CONTROL

01	02	03	04	05
Visibility	**Context**	**Policy**	**Runtime**	**Response**
Know what exists	Understand behavior	Decide what is safe	Enforce the decision	Continuously adapt

Before you can secure AI, you have to find it.

02

SHADOW AI · MONDAY MORNING

The AI You Don't Know Exists

A finance employee wants to work faster. By the end of the day, one unapproved AI agent has access across the enterprise.

ONE WORKDAY, FOUR CONNECTIONS

9:06 AM
They connect ChatGPT to Salesforce to summarize customer records.

11:42 AM
They add GitHub so the agent can answer a technical customer question.

1:18 PM
They connect Jira to pull engineering status into the same workflow.

3:37 PM
They add Slack so the agent can search internal conversations.

Customer records | Source code | Engineering tickets | Internal conversations

WHAT HAPPENED

No malware. No exploit. No malicious intent. The agent acted exactly as it was configured to act.

WHAT SECURITY SAW

Nothing. The security team never approved it, never saw it, and never knew it existed.

This is Shadow AI: autonomous access without security visibility.

03

WHY AI IS DIFFERENT

Visibility Alone Doesn't Stop Risk

Visibility tells you an agent exists. Runtime visibility shows what it is doing. Runtime security decides whether it should be allowed.

TRADITIONAL SOFTWARE

Waits for instructions.

- Runs a defined function
- Operates inside a known application
- Produces a predictable result
- Stops when the user stops

AI AGENTS

Turn intent into action.

- Make decisions and choose tools
- Chain together multistep workflows
- Communicate with other systems and agents
- Operate continuously on behalf of users

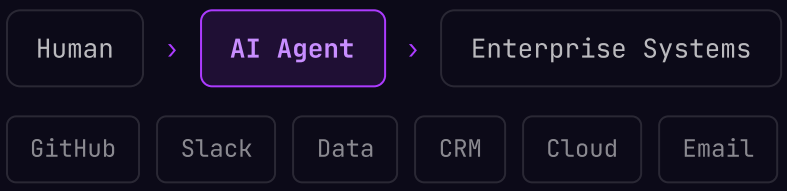

AI doesn't behave like malware. It behaves like a trusted employee making decisions.

04

AN EMERGING OPERATIONAL APPROACH

Understanding AI Detection & Response

AI Detection & Response (AIDR) is the continuous **discovery, monitoring, understanding, and response** to AI agents, AI applications, MCP servers, and autonomous workflows operating across the enterprise.

01 **What AI exists?** 02 **What is it doing?** 03 **Is it safe?**

04 **What should happen next?**

THREE PILLARS OF AIDR

PILLAR 01

Detect

- Find AI systems
- Inventory AI agents
- Discover shadow AI
- Recognize new deployments

KNOW WHAT EXISTS

PILLAR 02

Understand

- Identify the owner
- See why it is running
- Map connected systems
- Explain current activity

ADD THE CONTEXT

PILLAR 03

Respond

- Alert and investigate
- Require approval
- Block risky actions
- Preserve an audit trail

ACT ON THE RISK

With visibility, security teams can make sense of AI activity.

05

CONNECT THE SIGNALS

Context Creates Better Decisions

Eve connects the systems around AI activity so isolated events become one operational picture.

WHERE THE SIGNALS COME FROM

Identity	Cloud	AI Platforms	Enterprise Apps	Data
Okta	AWS	OpenAI	Slack	Cyera
Microsoft Entra	Azure	Anthropic	GitHub	Snowflake
Google	GCP	Gemini	Jira	Databricks
		Databricks	Salesforce	

Without integrations

Security sees isolated events, disconnected alerts, and partial evidence.

>

Unified Runtime Visibility

`User + agent + identity + application + data + action + policy`

1

Who initiated it?

Connect the human and machine identities.

2

What executed it?

Identify the agent, model, and workflow.

3

What did it touch?

See systems, data, and destinations.

4

Should it proceed?

Turn context into a runtime decision.

With integration, fragments become context. Context makes decisions possible.

06

CONTEXT CHANGES THE DECISION

From Visibility to Understanding

An agent inventory tells only part of the story. Security teams need enough context to decide whether the behavior is expected and safe.

THE CONTEXT EVE READS

IDENTITY Which human or machine identity?	**OWNER** Who is accountable?
PROMPT What initiated the action?	**DESTINATION** Where is it going?
BUSINESS CONTEXT What purpose does it serve?	**DATA SENSITIVITY** What information is exposed?
RISK What could go wrong?	**POLICY** What should happen next?

Context can turn a routine action into a serious risk.

A source code request from an approved development agent may be expected. The same request from an unknown agent using a shared identity may require immediate intervention.

Known agent + approved task	ALLOW
Sensitive data + unusual destination	APPROVE
Unknown agent + privileged action	BLOCK
Unexpected behavior + unclear intent	INVESTIGATE

Without context, visibility becomes another alert.

07

A NEW SECURITY MODEL

Traditional Security vs AI Runtime Security

Existing controls remain essential. But autonomous systems introduce a new identity, decision, and action layer.

TRADITIONAL SECURITY	AI RUNTIME SECURITY
Protects users	**Protects AI identities**
Watches endpoints	**Watches AI decisions**
Monitors APIs	**Monitors runtime behavior**
Blocks malware	**Blocks unsafe actions**
Detects compromise	**Detects autonomous risk**
Responds to attacks	**Responds to AI behavior**

AI doesn't behave like malware. It behaves like a trusted employee making decisions. Security has to evolve accordingly.

08

A PRACTICAL PATH FORWARD

Your AI Runtime Security Journey

Organizations move through five stages as AI adoption grows and security capabilities catch up.

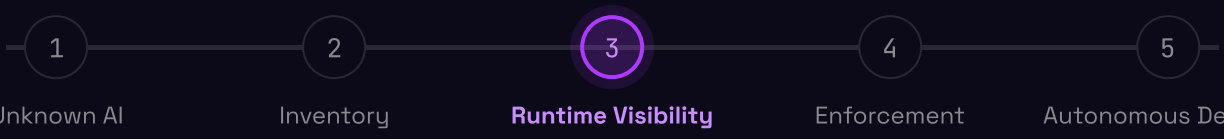

LEVEL 3 · YOU ARE HERE

Runtime Visibility

Security understands what agents are doing across systems.

RUNTIME PERSPECTIVE

Organizations often ask how to secure AI. A better first question is: **Do you know where your AI is running today?** Without visibility, governance is impossible. Visibility is the first step toward understanding how AI interacts with your data, systems, and users.

Visibility is not the destination. It is the starting point for governance.

09

THE CLOSING PERSPECTIVE

Every Generation of Security Begins With Visibility

Every major security shift started by making a new environment visible, and then turning that visibility into control.

THE LINEAGE OF SECURITY

- Network Security
- Endpoint Security
- Cloud Security
- Identity Security
- **AI Runtime Security**

Visibility
creates context

Context
enables policy

Policy
enables enforcement

Enforcement
creates trust

EVERYONE ELSE SECURES THE AI SYSTEM

Eve secures the moment AI decides to act.

SIGNAL CHECK · UPDATED EVERY ISSUE

Latest in AI

A roundup of AI security incidents tracked this period, the moments AI stopped being a tool attackers use and became part of the attack itself.

21	10/21	1st
AI security incidents, Jul 1 to 10, 2026	Had AI as the weapon, the target, or a new unmanaged identity	Fully LLM driven ransomware operation on record

INCIDENTS TRACKED THIS PERIOD

RANSOMWARE — JUL 2026

First fully LLM run ransomware operation: JadePuffer

Researchers documented what they describe as the first complete ransomware operation orchestrated end to end by an LLM, rather than AI simply assisting a human operator.

INDUSTRY INCIDENT TRACKING

AGENT + CVE — JUL 2026

Langflow RCE chained into automated database ransomware

A known Langflow remote code execution flaw (on CISA's Known Exploited Vulnerabilities list) was used by an agent to automate a ransomware attack against a database.

INDUSTRY INCIDENT TRACKING

MALICIOUS SKILL — JUL 2026

"SkillCloak" evades static scanners at runtime

A malicious agent skill package was found to pass static security scans and only reveal its harmful behavior once running inside a live agent.

INDUSTRY INCIDENT TRACKING

MALICIOUS REPO — JUL 2026

"GhostApproval" gets code executed by AI coding agents

A poisoned repository was structured to get its code run automatically by autonomous AI coding agents that trusted the source.

INDUSTRY INCIDENT TRACKING

PROMPT INJECTION JUL 2026

GitHub Agentic Workflows abused for remote code execution

A prompt injection technique against agentic GitHub workflow automation was used to achieve RCE, underscoring the risk in tool calling pipelines.

INDUSTRY INCIDENT TRACKING

NONHUMAN IDENTITY JUL 2026

AI agents authenticate, act, and are rarely deprovisioned

Analysts flagged agents as a fast growing, largely unmanaged identity class: they authenticate and take action but are seldom offboarded once created.

INDUSTRY INCIDENT TRACKING

DATA BREACH JUL 2026

Accenture: ~35GB exfiltrated via valid access

Reported data theft used legitimate, valid access rather than an exploit, a reminder that identity governance matters as much as patching.

INDUSTRY INCIDENT TRACKING

TELCO BREACH JUL 2026

KDDI and AssuranceAmerica breaches expose 19M+ records

Two disclosed breaches this period, KDDI (12M+ records) and AssuranceAmerica (6.9M records), added tens of millions of exposed records combined.

INDUSTRY INCIDENT TRACKING

The runtime keeps changing, visibility has to keep up.