CRANIUM

2026 EDITION

CRANIUM AI RESEARCH

The AI Supply Chain Risk Report.

Every model, agent and vendor you didn't build is now part of your attack surface. Third-party AI, AI-BOMs and agent supply chains: what 26 published sources show — and a five-move playbook to prove you're in control.

5 findings | 26 sources | 1 field playbook | 10-question self-assessment

2026

00 · EXECUTIVE SUMMARY

You're running more AI than you can see.

Most of the AI inside a modern enterprise was not built there. It arrives as open-weight models pulled from public hubs, as agents wired to tools through MCP servers, as features switched on inside SaaS you already bought, and as APIs from a handful of model providers. Together, these form the AI supply chain — and it is now one of the fastest-growing parts of the attack surface.

This report brings together the most credible public research of the past 20 months — from Verizon, IBM, Gartner, JFrog, OWASP, the G7 cyber agencies and regulators on both sides of the Atlantic — into one picture of where the AI supply chain is exposed, and what leading teams are doing about it.

Five findings

01 **The inventory is fiction.**
45% of employees now regularly use AI on corporate devices, up from 15% a year earlier.[1] Most AI arrives before anyone records it.

02 **Your vendor's AI is your attack surface.**
Third parties were involved in 48% of all breaches, up 60% in a year[1] — and AI is now shipping inside the products you already approved.

03 **Open model hubs are a live attack channel.**
Researchers counted 495 malicious AI models on Hugging Face in JFrog's 2026 study.[6] Scanning helps; attackers adapt.

04 **Agents widened the blast radius.**
Agents, MCP servers and skills marketplaces are the newest software supply chain — and they are drawing the oldest supply-chain attacks.

05 **Policy is table stakes. Evidence is the test.**
83% of S&P 500 companies now disclose AI as a risk.[17] Regulators, boards and customers want proof, not intent.

IBM X-FORCE, 2026[3]

> “Similar security incident rates were observed across both open-source and third-party vendor deployments.”

Buying the model doesn't outsource the risk.

ABOUT THIS REPORT

No new survey and no Cranium AI customer data. A synthesis of 26 public sources published February 2025 to October 2026. Vendor-funded research is marked [V] on the Sources page.

2026 AT A GLANCE

The AI supply chain, by the numbers.

48%
of all breaches involved a third party, up 60% year over year
VERIZON DBIR 2026 [1]

45%
of employees are regular AI users on corporate devices, up from 15%
VERIZON DBIR 2026 [1]

43%
of breaches involved shadow AI, up from 20% the year before
IBM COST OF A DATA BREACH 2026 [2, 26]

92%
of organizations with an AI-related breach lacked proper AI access controls
IBM X-FORCE 2026 [3]

150K+
agents predicted per global Fortune 500 enterprise by 2028, from fewer than 15 in 2025
GARTNER PREDICTION, APR 2026 [5]

72%
of financial institutions are only partially aware which vendors use AI
NCONTRACTS 2026 [8] [V]

495
malicious AI models found on Hugging Face
JFROG 2026 [6] [V]

250
poisoned documents backdoored every model tested, 600M to 13B parameters
ANTHROPIC · UK AISI · TURING [9]

83%
of S&P 500 companies disclose AI as a risk, up from 12% in 2023
THE CONFERENCE BOARD 2026 [17]

Figures as published by each source. 48% covers all breaches, not only AI-related ones. [V] marks vendor-funded or vendor-telemetry research. Full citations on page 14.

01 · DEFINITION

It's not just model weights.

Ask five vendors to define the AI supply chain and you'll get five answers. Model-security tools mean the files on Hugging Face. Third-party risk platforms mean the vendor questionnaire. Agent-security tools mean MCP servers. Each is right, and each is incomplete — attackers don't respect those boundaries, and neither do regulators.

DEFINITION

The AI supply chain is every model, agent, dataset and vendor that reaches your business — built, bought, downloaded or inherited — plus the code, tools and infrastructure that connect them.

AI supply chain risk management is the practice of inventorying, testing and continuously evidencing the AI components and suppliers an organization depends on, including AI it did not build.

FIVE LAYERS OF THE AI SUPPLY CHAIN

01 **Models**
Proprietary APIs, open-weight models, fine-tunes and adapters.

02 **Data**
Training, fine-tuning and retrieval data — and the prompts that shape behavior.

03 **Agents & tools**
Agent frameworks, MCP servers, plug-ins and skills.

04 **Vendors**
AI embedded in the SaaS, platforms and services you already buy.

05 **Code & infrastructure**
The repos, packages and cloud services that wire it all together.

BUILT BOUGHT DOWNLOADED INHERITED

What goes in an AI-BOM

An AI bill of materials is a machine-readable inventory of the models, datasets, frameworks, tools and infrastructure inside an AI system. In May 2026, G7 cyber agencies — including CISA, BSI, ANSSI and the UK NCSC — published voluntary minimum elements in seven clusters:[19]

- Metadata
- System-level properties
- Models
- Dataset properties
- Key performance indicators
- Infrastructure
- Security properties

The law already sees it

For high-risk AI systems, Article 25(4) of the EU AI Act requires providers and the third parties that supply “an AI system, tools, services, components, or processes” to agree in writing on the information, technical access and assistance needed to comply. For most high-risk systems, those obligations apply from December 2027.[12,15]



01

FINDING 01 · DISCOVER

The AI you can't see is the AI you can't defend.

Every governance program starts with an inventory. In most enterprises, that inventory describes the AI someone remembered to register — not the AI that is actually running.

Employees moved first. Verizon's 2026 DBIR found 45% of employees are now regular AI users on corporate devices, authorized or not, up from 15% a year earlier — and 67% reach AI services through non-corporate accounts.[1] Shadow AI is now the third most common non-malicious insider action in Verizon's data-loss data, and source code is the data type most often submitted to external AI.[1]

Builders are moving faster. Gartner predicts the average global Fortune 500 enterprise will have over 150,000 agents in use by 2028, up from fewer than 15 in 2025 — and only 13% of organizations think they have the right AI agent governance in place.[5]

The cost shows up in the breach data. In IBM's 2026 study, shadow AI was involved in 43% of breaches, up from 20%, and those incidents cost $5.39 million on average.[2,26]

> **"An inventory that depends on people remembering to register AI is a list of good intentions."**

THE EVIDENCE

45%
of employees regularly use AI on corporate devices (from 15%) [1]

67%
reach AI through non-corporate accounts on corporate devices [1]

13%
think they have the right AI agent governance in place [5]

$5.39M
average cost of breaches involving shadow AI [2, 26]

WHAT GOOD LOOKS LIKE

A living system of record — not a spreadsheet — that finds AI in source code and in the security telemetry you already collect, and generates an AI bill of materials for every system automatically.



02

FINDING 02 · THIRD-PARTY AI

Your vendor's AI is your attack surface.

Third parties were already the soft spot. Verizon's 2026 DBIR found breaches with third-party involvement rose 60% in a year, to 48% of all breaches.[1] Now AI is arriving inside those third parties' products, often as a feature switched on after the contract was signed.

That AI rarely triggers a new review. A vendor that passed assessment last year can add a model, an agent or an AI subprocessor this quarter. In Ncontracts' 2026 survey of financial institutions, 72% said they are only partially aware of which vendors use AI, and 9% haven't assessed it at all.[8]

CASE · AUGUST 2025

A connected AI chat app became the breach path.

Attackers used stolen OAuth tokens for Salesloft's Drift AI chat agent to export data from connected Salesforce instances. Google's Threat Intelligence Group said it was aware of over 700 potentially impacted organizations; Salesforce described a small number of affected customers.[20]

Five questions for every AI vendor

1. Which models, agents and AI subprocessors touch our data?
2. Is our data used to train or fine-tune any model?
3. How will you notify us before AI is added or changed?
4. Can you provide an AI-BOM in CycloneDX or SPDX?
5. What adversarial testing has your AI passed, and when?

THE EVIDENCE

48%

of all breaches involved a third party, up 60% [1]

72%

of financial institutions only partially aware which vendors use AI [8] [V]

700+

organizations potentially impacted through one AI chat-app integration [20]

WHAT GOOD LOOKS LIKE

Vendor AI that is disclosed, not assumed: a living transparency record for each supplier, requested and refreshed on a schedule you set — plus scoped, rotated tokens for every AI integration.



03

FINDING 03 · OPEN MODELS & DATA

Open model hubs are a distribution channel — for attackers too.

Open-weight models gave every team a head start — and gave attackers the same distribution channel that package registries gave them a decade ago.

JFrog's 2026 research counted 495 malicious AI models on Hugging Face and 177,000 new malicious packages across registries in the past year. It also found that while 97% of organizations claim certified model governance, 53% self-host models from sources where malicious payloads have been detected.[6]

Scanning helps, but attackers adapt. In February 2025, ReversingLabs found two malicious models on Hugging Face that used deliberately corrupted pickle files — a technique it named "nullifAI" — to slip past detection.[10]

The data layer is fragile too. Anthropic, the UK AI Security Institute and the Alan Turing Institute found that as few as 250 malicious documents implanted a backdoor in every model they tested, from 600 million to 13 billion parameters; the number did not grow with model size.[9] And AI now writes the supply chain: at least 5.2% of packages suggested by commercial models — and 21.7% from open-source models — did not exist, leaving names for attackers to register.[11]

> **"Scanning at download proves what you downloaded. It doesn't prove what is running today."**

THE EVIDENCE

495
malicious AI models on Hugging Face [6] [V]

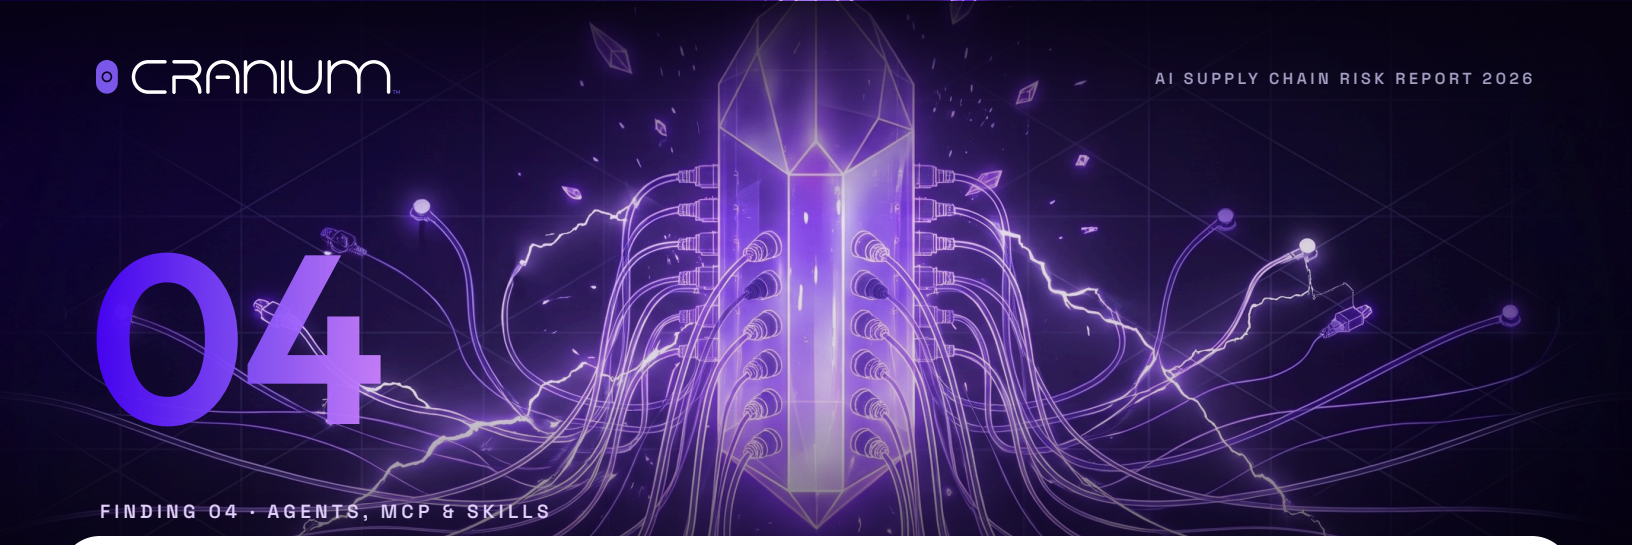
250
poisoned documents backdoored every model tested [9]

21.7%
of packages suggested by open-source models did not exist [11]

53%
self-host models from sources where malicious payloads were detected [6] [V]

WHAT GOOD LOOKS LIKE

Know which model, version and dependencies every system uses. Prefer allowlisted sources, safetensors and pinned hashes; test models adversarially before they ship and again whenever they change.



04

FINDING 04 · AGENTS, MCP & SKILLS

Every agent ships with a supply chain of its own.

An agent is a model plus tools, permissions and memory. Each tool connection — increasingly an MCP server or a downloadable skill — is a third-party component with the power to act. OWASP made it official in December 2025: its first Top 10 for Agentic Applications lists Agentic Supply Chain Vulnerabilities as ASI04.[13]

The warning signs arrived on schedule. In mid-2025, researchers disclosed a critical flaw (CVE-2025-6514, CVSS 9.6) in mcp-remote, a package used to connect AI clients to remote MCP servers and downloaded more than 437,000 times.[21] In September 2025, a malicious npm package impersonating Postmark's MCP server published 15 clean versions to build trust before silently copying outgoing email to an attacker; Koi Security called it the first malicious MCP server found in the wild.[22] In February 2026, an audit of 2,857 skills on one agent marketplace found 341 malicious ones.[23]

Adoption is running ahead of oversight. Microsoft's own telemetry shows more than 80% of Fortune 500 companies use active AI agents built with its low-code tools, and 29% of employees have turned to unsanctioned agents, according to a survey of security professionals.[14] JFrog found 18% of organizations have zero governance over their IDEs or MCP servers.[6]

> **"Every tool an agent can call is a supplier you may never have reviewed."**

THE EVIDENCE

80%+
of Fortune 500 use active agents built with Microsoft low-code tools [14] [V]

18%
have zero governance over IDEs or MCP servers [6] [V]

341
malicious skills found in one marketplace audit of 2,857 [23]

29%
of employees have turned to unsanctioned agents [14] [V]

WHAT GOOD LOOKS LIKE

An inventory of every agent and the frameworks, MCP servers and tools declared in its code — taken before deployment — plus runtime visibility into each agent's model traffic and the tool calls it makes. Scope every agent's tools to least privilege.

FINDING 05 · REGULATION & ASSURANCE

Policy is table stakes.
Evidence is the test.

Boards, regulators and customers have moved from asking whether an AI policy exists to asking for evidence that it is followed.

83% of S&P 500 companies now disclose AI as a risk, up from 12% in 2023.[17] Among the Fortune 100, 36% list AI as a separate 10-K risk factor.[25] And in a Vanta-commissioned survey of 2,500 leaders, 77% said stakeholders demand verified proof of compliance, up from 65% in 2024.[18]

U.S. banks face the opposite problem: no deadline at all. The April 2026 interagency model risk guidance placed generative and agentic AI outside its scope, with a separate request for information to follow.[16] Until then, the burden of proof sits with each institution's own program.

A delay is not a reprieve. The EU's high-risk deadlines moved; the evidence needed to meet them didn't. Organizations that use the extra time to build their AI-BOMs and vendor records will be ready to show their evidence when it is asked for.

Contract clause starter

Supplier shall (a) disclose all AI models, agents and AI subprocessors used to deliver the services; (b) give at least 30 days' written notice before adding or materially changing them; (c) not use Customer Data to train or fine-tune models without written consent; (d) provide an AI bill of materials in CycloneDX or SPDX format on request; and (e) notify Customer of any AI-related security incident within 72 hours.

A starting point for discussion with counsel, not legal advice.

THE REGULATORY CLOCK

AUG 2, 2025
EU AI Act obligations for general-purpose AI model providers apply.[15]

APR 17, 2026
U.S. banking agencies replace SR 11-7 model risk guidance — and place generative and agentic AI outside its scope.[16]

MAY 12, 2026
G7 cyber agencies publish minimum elements for an SBOM for AI.[19]

AUG 2, 2026
EU AI Act transparency duties (Art. 50) and the Commission's GPAI enforcement powers apply.[15]

JAN 1, 2027
Colorado's replacement AI law (SB 26-189) applies to consequential decisions; enforcement timing is subject to a pending federal ruling.[24]

DEC 2, 2027
EU AI Act high-risk obligations for Annex III systems apply, after the Digital Omnibus delay.[15]

AUG 2, 2028
EU AI Act high-risk obligations for AI in regulated products (Annex I) apply.[15]

EU dates reflect Regulation (EU) 2026/1744 (Digital Omnibus on AI), in force July 27, 2026.

83%
of S&P 500 companies disclose AI as a risk [17]

36%
of the Fortune 100 list AI as a separate 10-K risk factor [25]

77%
say stakeholders demand verified proof of compliance [18] [V]

06 · MARKET ANALYSIS

Every tool covers a slice.
Few cover the AI you buy.

Security teams have answered each new AI risk with a new tool. Our review of public documentation from more than 20 vendors found seven broad approaches. Each covers part of the AI supply chain well; the gap is a single record that follows one model, agent or vendor from discovery to proof — including the AI that lives inside your suppliers.

APPROACH	DISCOVER	OBSERVE	GOVERN	SECURE	PROVE	VENDOR AI
Model & ML supply-chain scanners Inspect model artifacts and packages	◐	◐	○	●	◐	○
Runtime guardrails & AI gateways Filter prompts and responses in flight	◐	●	◐	●	◐	○
Agent-security platforms Inventory and police agents and tools	●	●	◐	●	◐	○
Security platform suites AI modules inside broad security platforms	●	●	◐	●	◐	○
AI governance platforms Registries, approvals, policy mapping	◐	◐	●	◐	●	◐
Third-party risk platforms Questionnaires and vendor attestations	◐	◐	◐	○	◐	◐
Cloud, data & identity security Posture across cloud, data and identities	●	◐	◐	◐	◐	○

● CORE ◐ PARTIAL OR VARIES ○ LIMITED

THE TAKEAWAY

The gap isn't a missing stage. It's a missing record.

Security tools see the AI you run. Governance tools hold the policy. Third-party risk tools hold the vendor. Few join all three into one evidence trail per model, agent and vendor.

Six questions to ask any AI supply chain security platform

1. Does it find AI in source code, cloud and security telemetry — not just a registry?
2. Does it cover the AI you buy, not only the AI you build?
3. Can it export an AI-BOM in CycloneDX and SPDX?
4. Do red-team results land on the same record as the inventory?
5. Does it map to the EU AI Act, NIST AI RMF and ISO/IEC 42001 — not just threat lists?
6. Can it produce evidence for one model, agent or vendor on demand?

Ratings are Cranium AI's assessment of typical capabilities per category, based on public vendor documentation reviewed in October 2026 — not product testing. Individual products vary. Cranium AI sells into this market.

07 · FIELD PLAYBOOK

Five moves to secure the AI supply chain.

Run them in order — you can't govern what you haven't observed, or prove what you haven't secured. Then run them again.

01 DISCOVER
Build the system of record.

- Generate an AI-BOM from source code for every AI system: models, datasets, frameworks, infrastructure and system prompts.
- Inventory agents with the MCP servers and tools declared in their code.
- Pull shadow-AI signals from security telemetry you already collect.
- List every vendor that uses AI on your data — and how.

02 OBSERVE
See what AI actually does.

- Monitor prompts and responses in production, not just at approval.
- Record the tool calls agents make.
- Watch for drift between what was approved and what runs.

03 GOVERN
Turn policy into rules that run.

- Map each system to the NIST AI RMF, the EU AI Act and ISO/IEC 42001.
- Write AI clauses into supplier contracts (page 9).
- Set policy per use case, not one rule for all AI.

04 SECURE
Test before you trust. Defend while it runs.

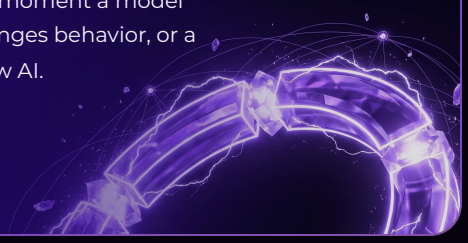
- Red-team models and agents before release, mapped to OWASP and MITRE ATLAS.
- Enforce runtime policies on sensitive data, secrets, jailbreaks, instruction overrides and direct injection attempts.
- Track CVEs in AI components; scope agent tools to least privilege.

05 PROVE
Make evidence a byproduct, not a project.

- Keep a living transparency record for every internal system and vendor.
- Export AI-BOMs in CycloneDX and SPDX.
- Answer the board, the regulator and the customer questionnaire on demand.

CONTINUOUSLY
The loop, not a sixth step.

Rerun the loop the moment a model ships, an agent changes behavior, or a vendor turns on new AI.



08 · SELF-ASSESSMENT

Where does your AI supply chain stand?

Answer yes or no. Score one point for each yes, then find your level. Be strict: “we're working on it” is a no.

- [ ] 01 Can you list every AI model in production today, with its version?
- [ ] 02 Do you know which agents run in your environment and which tools and MCP servers they use?
- [ ] 03 Can you see AI use that never went through procurement?
- [ ] 04 Do you know which of your vendors use AI on your data?
- [ ] 05 Do your contracts require vendors to notify you before adding or changing AI?
- [ ] 06 Are AI systems red-teamed before release and after material change?
- [ ] 07 Do runtime policies inspect prompts and responses for sensitive data and attacks?
- [ ] 08 Is each AI system mapped to the NIST AI RMF, the EU AI Act or ISO/IEC 42001?
- [ ] 09 Can you produce an AI-BOM in a standard format on request?
- [ ] 10 Could you assemble an evidence pack for one AI system within five business days?

YOUR SCORE

0–3	Level 1 · Unaware
4–6	Level 2 · Inventoried
7–8	Level 3 · Governed
9–10	Level 4 · Provable

Four metrics to track

- Share of AI systems with a current AI-BOM
- Share of AI-using vendors with a current disclosure
- Days from AI change to updated record
- Time to assemble an evidence pack

YOUR NEXT MOVE

LEVEL 1

Build the inventory first: questions 1–4. You can't govern what you haven't found.

LEVEL 2

Automate it, and put AI notice into every supplier contract: questions 5 and 9.

LEVEL 3

Close the evidence loop so proof is on demand: question 10.

1 Unaware
AI is tracked in spreadsheets, if at all. Vendor AI is unknown.

2 Inventoried
A registry exists, but it's manual and goes stale between reviews.

3 Governed
Inventory is automated, policies map to frameworks and testing happens before release.

4 Provable
Every model, agent and vendor is discovered, observed, governed, secured and proved — continuously, with evidence on demand.

Scored 6 or less? Start with questions 1, 4 and 9.
See what's already running before you plan what to govern.

Run a free AI exposure scan →

HOW CRANIUM AI HELPS

The AI you build. The AI you buy.

One record.



We know what it's like to be accountable for AI you can't fully see — to be asked for certainty about a system that won't sit still.

Cranium AI is the AI security and governance platform built for the enterprise. Cranium ONE brings your models, agents and vendors into one AI Trust Loop — discover them, observe them, govern them, secure them and prove it, continuously.

01 DISCOVER

AI-BOMs generated automatically from GitHub, GitLab, Bitbucket and Azure DevOps repositories. AgentSensor™ inventories agent frameworks, MCP servers and tools from source code. A shadow-AI view from your connected SIEM.

02 OBSERVE

Guardian inspects each interaction between connected AI tools and the models they call, records the tool calls agents make, and explains every verdict with Trace™ — deterministic classifiers, not an LLM judging an LLM.

03 GOVERN

Policy profiles per use case, with compliance mapped to the EU AI Act, NIST AI RMF and ISO 42001 — and attestations backed by citations.

04 SECURE

Arena™ automated and manual red teaming mapped to MITRE ATLAS and OWASP. In Enforce mode, runtime profiles block or modify prompts and responses carrying PII, PHI, PCI, secrets, jailbreaks or instruction overrides.

05 PROVE

AI Cards™: living, exportable transparency reports for your AI systems and the vendors you request them from, with AI Trust Hubs. AI-BOMs export in CycloneDX and SPDX.

CONNECTS TO GitHub · GitLab · Bitbucket · Azure DevOps · OpenAI · Anthropic · AWS Bedrock
MAPS TO MITRE ATLAS · OWASP · EU AI Act · NIST AI RMF · ISO 42001 **EXPORTS** CycloneDX · SPDX

Works with the stack you already have | Your policy stays yours | One platform, one motion

Named a Gartner® Cool Vendor in AI Cybersecurity Governance, one of Fortune's Top 50 cybersecurity companies, and a TAG Top 5 AI security vendor. Named a Niche Player in the 2026 Gartner® Magic Quadrant™ for AI Governance Platforms.

See your own AI supply chain, discovered live.
cranium.ai

Get a Demo →

METHODOLOGY & SOURCES

How this report was built

This edition is a synthesis of published third-party research, regulatory texts and public incident reporting released between February 2025 and October 2026. Figures are reported as published by each source; survey sizes and sponsors are noted where relevant, and vendor-funded or vendor-telemetry research is marked [V]. Forecasts are labeled as predictions. The market analysis on page 10 reflects Cranium AI's review of public vendor documentation in October 2026, not product testing. No Cranium AI customer data was used. Research current as of October 29, 2026.

1 Verizon Business, *2026 Data Breach Investigations Report*, Executive Summary, May 19, 2026. verizon.com

2 IBM, "IBM Study: One in Four Malicious Breaches Are AI-Enabled, Costing Companies $6 Million on Average" (*Cost of a Data Breach Report 2026*), July 29, 2026. newsroom.ibm.com

3 IBM X-Force, "AI-powered adversaries and the enterprise risk challenge," July 29, 2026. ibm.com

4 IBM, *Cost of a Data Breach Report 2025*, press release, July 30, 2025. newsroom.ibm.com

5 Gartner, "Gartner Identifies Six Steps to Manage AI Agent Sprawl," press release, April 28, 2026. gartner.com

6 JFrog, *2026 Software Supply Chain Security State of the Union*, press release, May 20, 2026 (n=1,508). [V] s21.q4cdn.com

7 Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," press release, August 26, 2025. gartner.com

8 Ncontracts, *2026 State of Third-Party Risk Management Survey*, June 4, 2026. [V] ncontracts.com

9 Anthropic, UK AI Security Institute and The Alan Turing Institute, "A small number of samples can poison LLMs of any size," October 9, 2025. anthropic.com

10 ReversingLabs, "ReversingLabs Identifies Novel ML Malware Hosted on Leading Hugging Face AI Model Platform," February 6, 2025. reversinglabs.com

11 Spracklen et al., "We Have a Package for You! A Comprehensive Analysis of Package Hallucinations by Code Generating LLMs," USENIX Security Symposium 2025. arxiv.org

12 Regulation (EU) 2024/1689 (EU AI Act), Article 25, "Responsibilities Along the AI Value Chain." eur-lex.europa.eu

13 OWASP GenAI Security Project, *OWASP Top 10 for Agentic Applications for 2026*, December 9, 2025. genai.owasp.org

14 Microsoft Security Blog, "80% of Fortune 500 use active AI Agents: Observability, governance, and security shape the new frontier," February 10, 2026. [V] microsoft.com

15 European Commission, "Guidelines for providers of general-purpose AI models" digital-strategy.ec.europa.eu; Regulation (EU) 2026/1744 of 8 July 2026 (Digital Omnibus on AI) eur-lex.europa.eu; Gibson Dunn, "EU AI Act Omnibus Agreement — Postponed High-Risk Deadlines and Other Key Changes," May 27, 2026.

16 Davis Polk, "Visual memo: Key changes under the federal banking agencies' revised model risk management guidance," April 2026 (OCC, FDIC and Federal Reserve guidance issued April 17, 2026). davispolk.com

17 The Conference Board, *From Principles to Practice: Governing AI in the Corporation*, press release "In the S&P 500, Disclosure of AI Risks Surges from 12% to 83%," April 22, 2026 (S&P 500 disclosure data as of December 2025). conference-board.org

18 Vanta, *State of Trust Report* (3rd edition), cited in "The top 6 AI security trends for 2026," December 17, 2025. [V] vanta.com

19 G7 Cybersecurity Working Group, *Software Bill of Materials (SBOM) for Artificial Intelligence – Minimum Elements*, published by BSI, May 12, 2026. bsi.bund.de

20 The Hacker News, "Salesloft OAuth Breach via Drift AI Chat Agent Exposes Salesforce Customer Data," August 27, 2025. thehackernews.com

21 The Hacker News, "Critical mcp-remote Vulnerability Enables Remote Code Execution," July 10, 2025 (JFrog Security Research disclosure, CVE-2025-6514). thehackernews.com

22 The Register, "postmark-mcp" malicious MCP server report, September 29, 2025 (Koi Security research; Koi describes it as the first malicious MCP server found in the wild). theregister.com

23 The Hacker News, "Researchers Find 341 Malicious ClawHub Skills," February 2, 2026 (Koi Security research). thehackernews.com

24 Jones Walker LLP, "The Colorado AI Law You Prepared For Is Gone, and the Rulemaking Comment Window Closes October 26," October 23, 2026. joneswalker.com

25 EY Center for Board Matters, "Cyber and AI Oversight Disclosures: What Companies Shared in 2025," Harvard Law School Forum on Corporate Governance, October 28, 2025. corpgov.law.harvard.edu

26 Baker Donelson, "Ten Takeaways from IBM's 2026 Cost of a Data Breach Report," 2026 (summary of IBM *Cost of a Data Breach Report 2026* shadow-AI findings). bakerdonelson.com

Gartner, Cool Vendors in AI Cybersecurity Governance, 2025. Gartner, Magic Quadrant for AI Governance Platforms, June 2026. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose. GARTNER is a registered trademark and service mark, and MAGIC QUADRANT and COOL VENDORS are registered trademarks, of Gartner, Inc. and/or its affiliates in the U.S. and internationally and are used herein with permission. All rights reserved. Gartner statistics cited are Gartner predictions from Gartner press releases. Cranium ONE, AgentSensor, Arena, AI Card and Trace are trademarks of Cranium AI. This report is for information only and is not legal advice.

CRANIUM AI

Adopt AI fast.

Prove it's safe.

Every model, agent and vendor — discovered, observed, governed, secured and proved. Continuously.

Get a Demo →

CRANIUM.AI



DISCOVER · OBSERVE · GOVERN · SECURE · PROVE — CONTINUOUSLY

CRANIUM™

Cranium AI is the AI security and governance platform built for the enterprise. Cranium ONE discovers, observes, governs, secures, and proves every model, agent, and vendor across the enterprise AI supply chain. © 2026 Cranium AI.