GC | CYBERSECURITY

STRATEGIC WHITE PAPER · 2026

The Autonomous Data Protection Imperative

Why enterprises must rethink data exfiltration risk in the age of AI and autonomous agents — and how to close the gaps between the controls they already own.

A briefing for CISOs and CIOs of global enterprises and government agencies

September 2026

EXECUTIVE SUMMARY

Enterprise data protection is entering a new era

Enterprises have invested heavily in classification, DLP, CASB, SASE and DSPM. Those investments remain essential — but the way sensitive data moves has changed faster than the architecture built to protect it.

Data now flows through cloud services, collaboration platforms, generative AI tools and, increasingly, autonomous agents. An employee can expose a contract in a single prompt. An approved agent can retrieve, summarize and transmit data with no human reviewing each step. A compromised identity can use entirely legitimate channels to move it out.

The challenge is no longer simply *identifying* sensitive data. It is knowing how it is used, where it is moving, who — or what — is moving it, and whether that is authorized; then acting before exposure becomes an incident. Each existing control was built around a particular channel, data state or question, so gaps now emerge *between* them — even when every product works exactly as designed.

> *The question is not whether your controls work. It is whether, together, they cover the pathways your data actually travels.*

This paper maps where those seams appear, offers an eight-dimension framework and a 0–100 score to measure them, and sets out how an autonomous protection layer — complementary by design — can be validated in your own environment.

$4.99M
record global average cost of a data breach, up 12% in a year[1]

43%
of breached organizations reported shadow-AI incidents, up from 20%[1]

223
GenAI data policy violations per organization every month — double the prior year[3]

48%
of breaches involved a third party, up from 30%[2]

THREE CONCLUSIONS FOR SECURITY LEADERSHIP

01
Reassess the architecture, not the products
Point-product scorecards will not reveal the seams between them.

02
Treat AI and agents as first-class pathways
They now carry regulated data, source code and IP at machine speed.

03
Prove prevention with evidence
Validate policy-driven prevention in a scenario-based proof of value.

01 · THE LANDSCAPE

The exfiltration problem has changed — and AI has outpaced the controls

Data exfiltration is the unauthorized movement of sensitive information beyond intended boundaries. Two of the six pathways below barely existed in risk registers three years ago; today they carry a rapidly growing share of enterprise data, at machine speed.

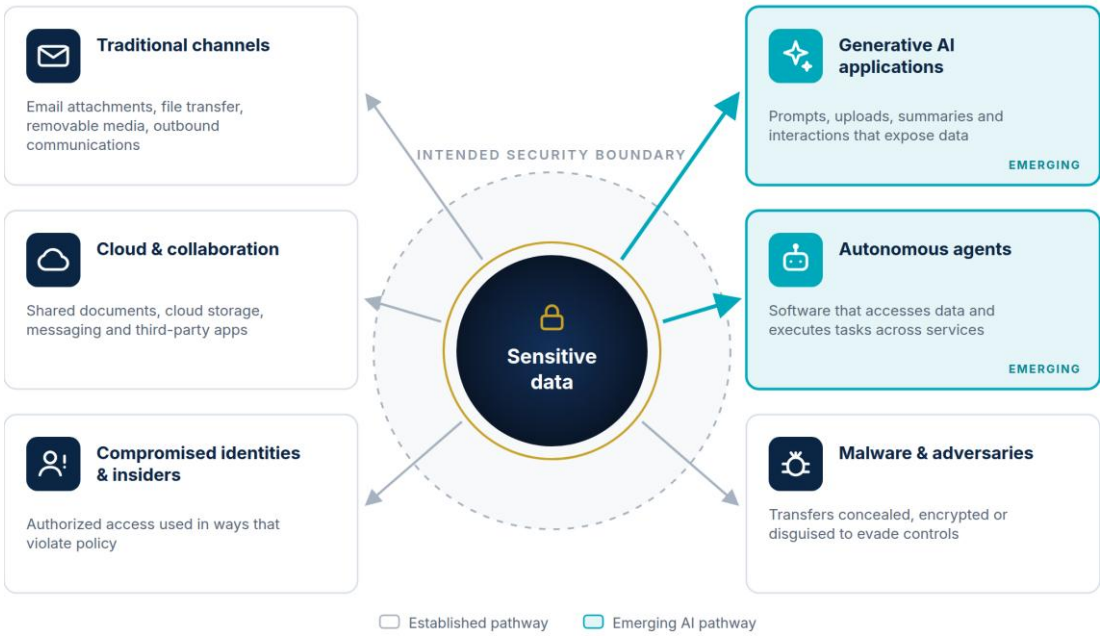


FIGURE 1 Established pathways remain; emerging AI pathways add new routes, new actors and machine-speed volumes.

6×
growth in prompts sent to SaaS AI apps in one year — 3,000 to 18,000 a month[3]

39.7%
of employee AI interactions involve sensitive data[7]

90%
of GenAI policy violations involve source code, regulated data or IP[3]

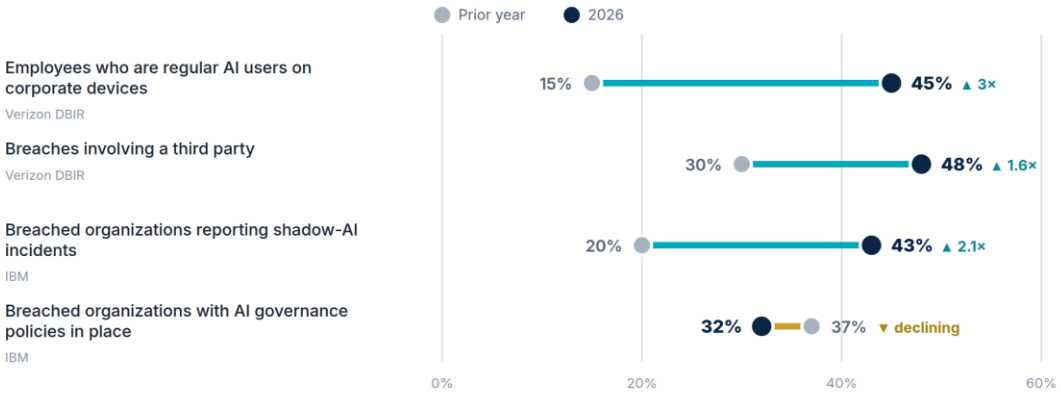


FIGURE 2 One year, four structural shifts: adoption and exposure rise while AI governance declines. Sources: Verizon 2026 DBIR[2]; IBM 2026.[1]

02 · THE ARCHITECTURE

Protection gaps emerge between controls — not within them

A control can be highly effective inside its design center — a channel, a data state, or a question such as "who has access?" — while offering little elsewhere. The pathways created by AI and automation increasingly fall between those design centers. The question is not whether these technologies are useful; it is whether the combined architecture covers the organization's actual data-movement pathways.

	Email & file transfer	Sanctioned SaaS & cloud	THE SEAMS — NO PRIMARY PREVENTIVE OWNER: GenAI prompts & uploads	Sanctioned AI agents	Shadow AI & rogue agents	Insider & compromised identity	Obfuscated / encrypted exfiltration
Data classification & labeling	◑	◑	○	○	○	○	○
Data loss prevention (DLP)	◑	◑	◑	○	○	◑	○
Cloud access security broker (CASB)	◑	●	◑	○	○	◑	○
SASE / security service edge	◑	●	◑	◑	◑	○	○
Data security posture mgmt (DSPM)	○	◑	○	◑	○	○	○
Identity & privileged access (IAM/PAM)	○	◑	○	◑	○	◑	○
SIEM / UEBA	◑	◑	○	○	○	◑	○
Primary preventive owners	**0**	**2**	**0**	**0**	**0**	**0**	**0**
GC Autonomous Data Protection — design focus	Uses existing signals		Autonomous detection & prevention across the seams				

● Primary design center — preventive control ◑ Partial, conditional or detect-only ○ Limited or indirect by design

FIGURE 3 Typical design-center coverage of established controls by pathway. Illustrative and category-level; actual coverage varies by vendor, configuration and deployment.

6 of 7 **pathways have no control whose primary design center is prevention there** — only sanctioned SaaS has one. Not because existing tools fail, but because no single tool was built to own them.

02 · COMPLEMENT, DON'T REPLACE

A context-and-decision layer that strengthens what you already own

Existing controls keep doing what they do best and contribute their signals. The autonomous layer unifies that context, applies policy consistently across pathways, enforces where others only observe — and records the evidence of every decision.

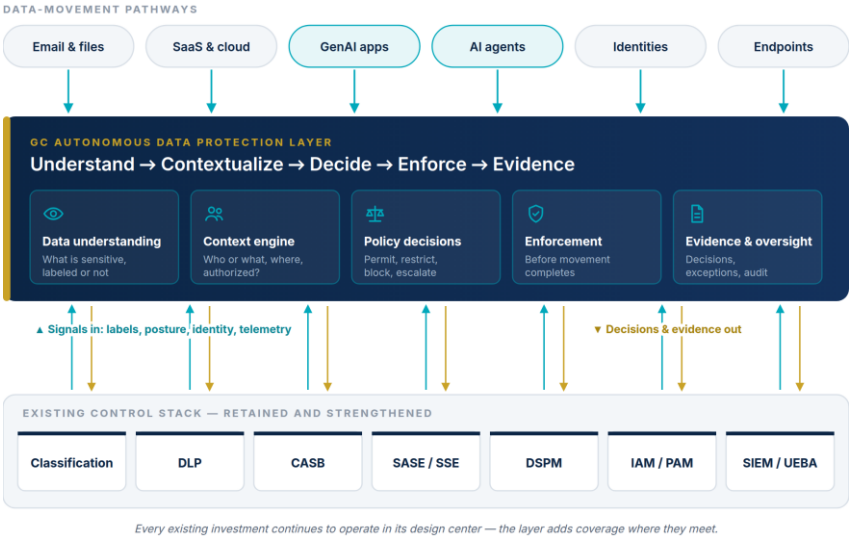


FIGURE 4 Existing controls contribute signals; the layer contributes cross-pathway decisions, enforcement and evidence.

Existing control	Where the seam typically appears	How autonomous protection complements it
Data classification & labeling	Labeling is rarely comprehensive, and malicious insiders can deliberately mis-classify data to bypass DLP controls. Labels rarely survive copy-paste, AI summarization or derived content.	Automatically identifies and classifies data in motion, with no requirement for manual labeling. Uses pre-labeling as guidance only.
Data loss prevention (DLP)	Lacks data exfiltration prevention controls. Blind to pathways outside its sensors; lacks identity and intent context; high false positives limit its usefulness.	Adds identity, destination and behavioral context; extends to AI and agent pathways. Critically, provides autonomous data exfiltration prevention controls.
CASB	Like DLP, lacks data exfiltration prevention controls. Often retrospective; thins for unsanctioned apps and personal accounts.	Evaluates movement as it happens, including to AI services outside the catalog. Critically, provides autonomous data exfiltration prevention controls.
SASE / SSE	Sees only what traverses it; agent-to-API and local flows can fall outside	Applies consistent data-centric policy where traffic-centric controls lack context
DSPM	Point-in-time posture; does not act on data in motion	Turns posture insight into runtime decisions about data movement and data protection controls
IAM / PAM	Says nothing about what an authorized identity does with data next	Monitors how authorized identities — human or agent — actually move data
SIEM / UEBA	Detection generally follows the event; data has often already left	Contributes prevention before transfer completes, and evidence to the CyberSOC with real-time reports and forensic audit trails

03 · THE NEW RISK SURFACE

AI and autonomous agents multiply the ways data can move

The risk is not limited to employees pasting data into chatbots. Every agent is a new identity with access to data and the ability to move it — and exposure depends on what it can reach and what governs it.

40%
of enterprise apps to embed AI agents by end-2026, from <5%[5]

45:1
non-human identities per human in the average enterprise[8]

67%
of corporate-device AI users rely on non-corporate accounts[2]

92%
of orgs with an AI incident lacked proper AI access controls[1]

3.1
Sensitive data submitted to AI
Prompts, attachments and uploads expose confidential material to external services.

3.2
Sanctioned AI agents
Approval of the agent is not approval of every transfer it makes.

3.3
Unsanctioned or unknown agents
Outside governance: unclear access, flows and accountability.

3.4
Workflows at machine speed
Decide what to automate, what to escalate, what to block.

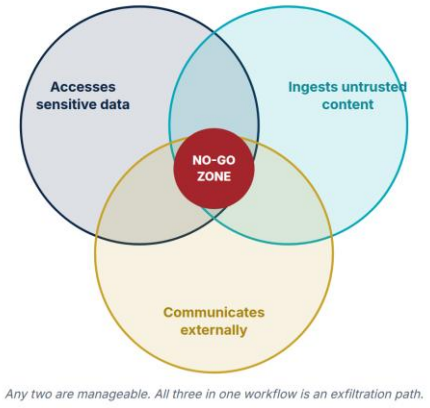


FIGURE 5 The agentic “no-go zone”. Concept: Gartner (2026).[6]

When an agent can read, ingest and send, risk compounds

Gartner flags agents that access sensitive data, ingest untrusted content and communicate externally in one workflow as “no-go zones”: a single injected instruction can turn an approved agent into an exfiltration channel.[6] The answer is not to remove humans, but to place their judgment where it matters:

- **Automate** routine, policy-conforming activity — every decision logged
- **Escalate** high-impact or ambiguous transfers for human approval
- **Block** clearly unauthorized movement before it completes

04 · THE CONSEQUENCES

AI and delay are driving the cost of exfiltration — detection alone is too late

Beyond direct cost, exfiltration brings operational disruption, regulatory exposure, IP loss, damaged trust and strategic disadvantage. Industry benchmarks show two consistent multipliers: AI involvement and time.

The cost ladder of a single breach

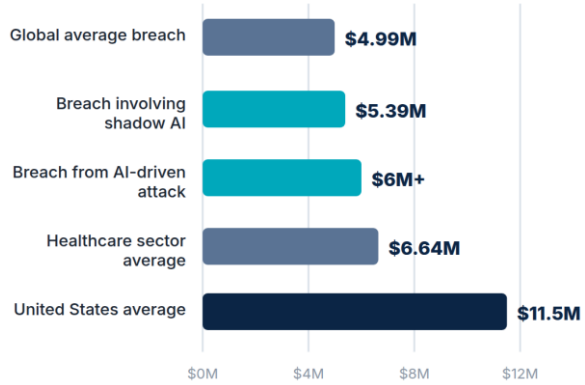


FIGURE 6 Average breach cost, US$M, 2026.[1]

Insider risk: time is money

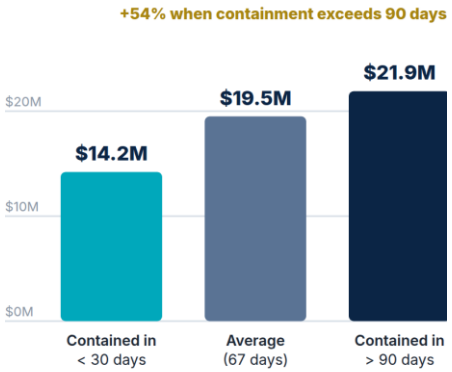


FIGURE 7 Annual insider-risk cost by time to contain, US$M.[4]

$19.5M
average annual insider-risk cost per organization[4]

247
days, the average breach lifecycle — the first rise in five years[1]

+56%
year-over-year increase in AI-driven attacks[1]

Why detection alone is not enough

Monitoring often identifies activity after data has left. Prevention stops the transfer before it completes — the difference between an alert and an incident that never happened.

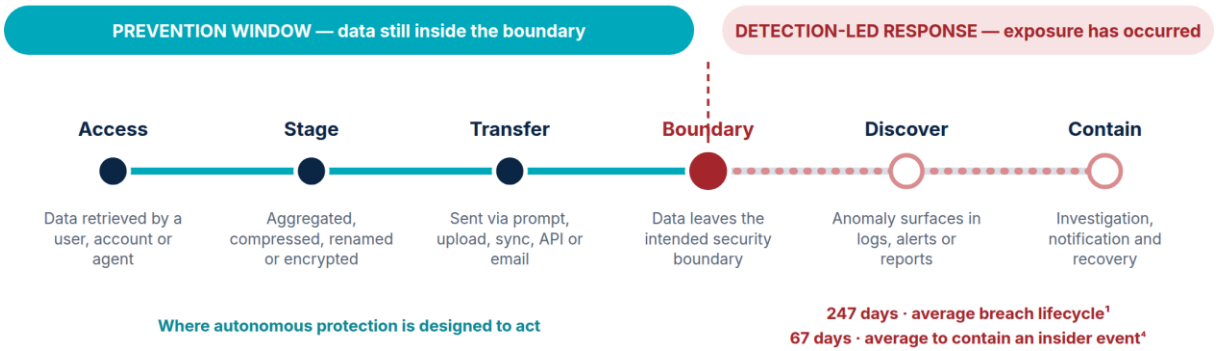


FIGURE 8 Prevention operates before data crosses the boundary; detection-led controls typically act after it.

05–06 · THE FRAMEWORK

Eight dimensions of risk. One 0–100 score.

GC recommends assessing eight dimensions that move the conversation from “which products do we own?” to “what can we actually see, decide and stop?” Each is scored 0–4 (higher means greater exposure) and weighted into a composite 0–100 score.



FIGURE 9 Exposure (navy), detection & response (teal), resilience (gold).

1. **Sensitive data exposure** What is held, where, and who or what can access it?
2. **Data movement & channel coverage** Which channels are monitored; which workflows fall outside?
3. **AI & agentic exposure** Can AI apps and agents be identified and governed, sanctioned or not?
4. **Insider & identity risk** Can authorized access be told apart from misuse or compromise?
5. **Detection & contextual understanding** Are identity, destination, behavior and authorization weighed?
6. **Prevention & autonomous response** Can transfers be stopped before data leaves, with human oversight?
7. **Evasion & adversarial resilience** Are encryption and obfuscation detected — and tested?
8. **Governance & operational readiness** Can control operation be demonstrated?

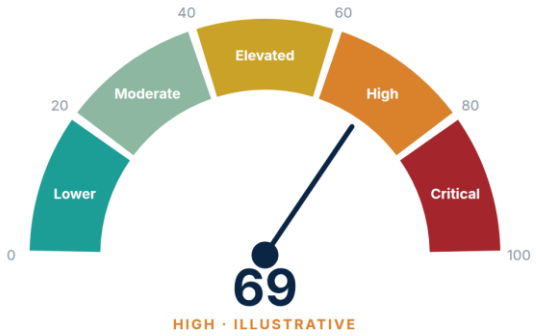


FIGURE 10 Composite score bands; the needle shows an illustrative profile scoring 69.

Score	Risk classification
0–19	Lower
20–39	Moderate
40–59	Elevated
60–79	High
80–100	Critical

Not a prediction of breach probability. The score identifies weaknesses, prioritizes validation and sets a baseline to report improvement — interpreted alongside industry, data sensitivity and threat environment.

07–09 · THE APPROACH

The case for autonomous data protection

Autonomous data protection uses context, policy and automated decision-making to identify and respond to data risks with reduced dependence on manual intervention. Effective protection connects five capabilities — and enforcement is the one detection-led architectures lack.

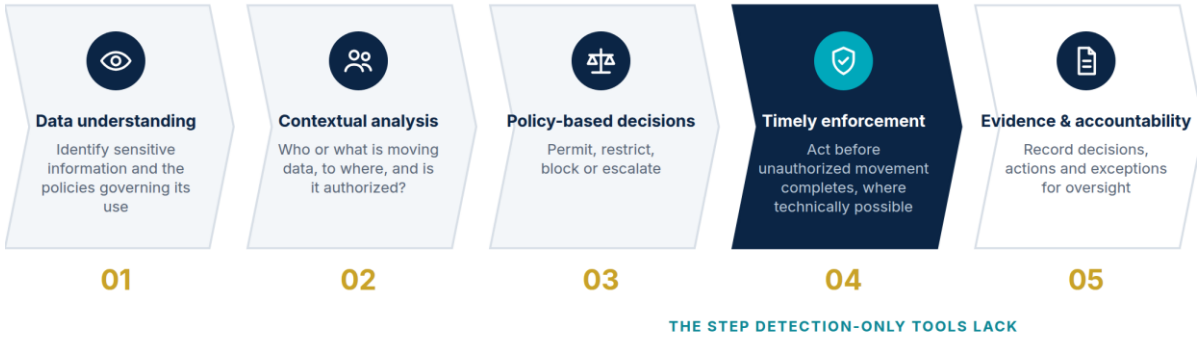


FIGURE 11 The protection chain. No single control should be assumed to prevent every exfiltration scenario.

Six capabilities any platform in this category must prove

Continuous visibility

Real-time automated data classification, without any human intervention, across all channels.

Automated data & actor awareness

Sensitive data, human or non-human identities, destinations and activity.

Automated policy-driven protection

Granular rules applied consistently across all actors and workflows.

Real-time automated enforcement

Accidental and malicious data leak prevention, GenAI guardrails, regulatory compliance.

Human oversight

Governance, escalation and approval where judgment matters.

Operational integration

Works alongside — and strengthens — the existing stack.

GC's Autonomous Data Protection Platform

Designed to protect sensitive information across modern enterprise workflows by understanding data and actor activity and supporting autonomous protection against unauthorized movement — working alongside, and drawing strength from, the controls already in place.

The strongest case for deployment is a measurable improvement in coverage, detection quality or prevention effectiveness — demonstrated in the customer's own environment.

FIVE QUESTIONS TO PUT TO ANY PLATFORM — INCLUDING OURS

1. Can it automatically identify and classify sensitive information moving through **key** channels?
2. Can it recognize suspicious or unauthorized activity — by humans or NHI, including **AI agents**?
3. Can it **prevent** sensitive data from leaving the intended boundary in real time?
4. Can it address the **data leak and exfiltration scenarios** our existing controls do not adequately cover?
5. Can it provide the **forensic evidence** and controls enterprise deployment requires?

10–11 · FROM ASSESSMENT TO DECISION

From risk assessment to a business case a CFO will trust



FIGURE 12 Map, assess, validate, decide: the four-phase path from assessment to deployment.

Nine critical use cases to validate first

Data exfiltration use case	Type	Actor	Channel	Risk
Mis-directed email with sensitive content	Accidental	Insider	Email	High
Accidental upload of sensitive unstructured content	Accidental	Insider	GenAI	Severe
Known AI agent accidentally leaks sensitive content	Accidental	Known AI agent	GenAI	Severe
Rogue AI agent exfiltrates sensitive content	Malicious	Rogue AI agent	GenAI	Critical
Malware exfiltration with obfuscation of content	Malicious	Malware	Email	Critical
Encryption bypass of sensitive content	Malicious	Insider	Email	Severe
Cloud draft exfiltration of sensitive content	Malicious	Insider	Cloud	Severe
Teams upload of sensitive unstructured content	Malicious	Insider	Teams	Severe
FTP bulk upload (>100 MB) of sensitive content	Malicious	Insider	FTP	Severe

Proof-of-value criteria: coverage · detection · prevention · latency · false-positive and false-negative rates · governance · integration.

Value lever	Evidence class
Exposure reduction	MEASURED IN PROOF OF VALUE
Protection coverage	MEASURED IN PROOF OF VALUE
Response effectiveness	MEASURED IN PROOF OF VALUE
Risk governance	MEASURED IN PROOF OF VALUE
Operational efficiency	MODELED ON CUSTOMER DATA
Integration cost	MODELED ON CUSTOMER DATA
Avoided loss	STATED ASSUMPTION

INDUSTRY BENCHMARK

$1.93M

lower breach cost with extensive security AI and automation; lifecycles 65 days shorter.[1] Context only — not a GC claim.

THE CORE EQUATION

Net value = (exposure closed × customer-specific impact) + efficiency gained – (deployment + operating + integration cost)

Protecting data in an autonomous world

Existing technologies remain important, but no single control guarantees protection across every pathway. The imperative is to assess coverage, identify material gaps, and validate whether the architecture can close them.

FIVE QUESTIONS FOR YOUR NEXT BOARD BRIEFING

Q1 Which of our data-movement pathways sit outside the effective coverage of any preventive control today?
Q2 Can we see — and stop — sensitive data flowing into AI tools and agents, sanctioned or not?
Q3 If a trusted identity began moving data at machine speed tonight, would we prevent it or read about it later?
Q4 Could we evidence to a regulator that our data controls operated as intended?
Q5 Which existing investments would deliver more value with better context?

Make sensitive data movement visible. Make unauthorized activity detectable. Make effective protection actionable.

NEXT STEP

Start with a Data Protection Risk Assessment

A 0–100 baseline across eight dimensions, a map of the seams between your existing controls, and a proof of value on the nine critical use cases — in your own environment.

ENDNOTES

1. IBM Security & Ponemon Institute, Cost of a Data Breach Report 2026.
2. Verizon, 2026 Data Breach Investigations Report (22,000+ confirmed breaches).
3. Netskope Threat Labs, Cloud and Threat Report: 2026.
4. Ponemon Institute, 2026 Cost of Insider Risks: Global Report (354 organizations).
5. Gartner, press release, 26 August 2025 (task-specific AI agents in enterprise apps).
6. Gartner, press release, 9 April 2026 (GenAI security incidents; agentic “no-go zones”).
7. Cyberhaven, 2026 AI Adoption & Risk Report.
8. Cloud Security Alliance, 2026 non-human identity survey, via Security Boulevard (July 2026).