Executive White Paper · 2026

AI Exposure Management

The Operating Model for Governed Enterprise AI

Discover every AI agent. Identify ownership. Map identities, tools, data, actions, spend, and exposure paths. Prove governance with evidence.

For CISOs, CIOs, CFOs, Security, IT, Risk, and Finance leaders

86%

Employers expecting AI to transform business by 2030

33%

Enterprise software applications expected to include agentic AI by 2028

15%

Day-to-day work decisions expected to be made autonomously by 2028

$670K

Incremental average breach cost linked to Shadow AI

Sources used in this report include World Economic Forum, Gartner, NIST, European Commission, and IBM research as reported by secondary analyses.

Executive Summary

AI is becoming enterprise infrastructure faster than enterprise control models are adapting.

1. AI is shifting from experimentation to infrastructure, creating a new enterprise control challenge.
2. Agentic systems introduce identity, connector, tool, data, action, and ownership risk that traditional Shadow AI controls may miss.
3. Governed adoption requires discovery, ownership, policy, spend visibility, and evidence-backed oversight.

Chart 1 · AI Transformation Pressure

86%

of employers expect AI and information-processing technologies to transform their business by 2030.

Source: World Economic Forum, Future of Jobs Report 2025.

Chart 2 · The Agentic Adoption Curve

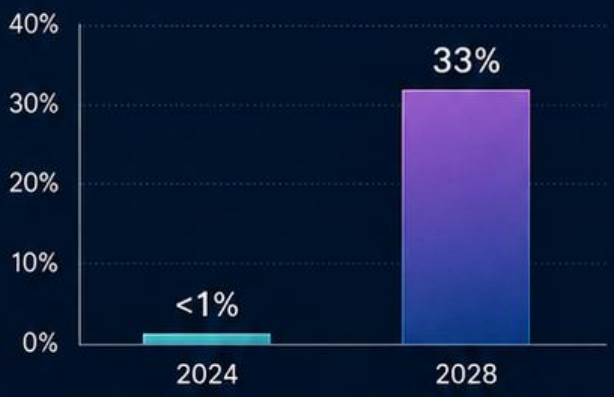

Enterprise software applications expected to include agentic AI.

Source: Gartner, June 2025.

What this means

AI governance can no longer be treated as an innovation side project. As agentic capabilities become embedded in enterprise applications, organizations need an operating model that can discover AI systems, assign ownership, map exposures, and prove control.

- Discover every AI system
- Assign clear ownership
- Map and manage exposures
- Gain visibility into AI spend
- Prove control with evidence and auditability

From Assistance to Autonomous Action

As agentic AI moves from answering questions to taking action, control requirements shift from visibility to governance.

Chart 3 · Autonomous Decisions

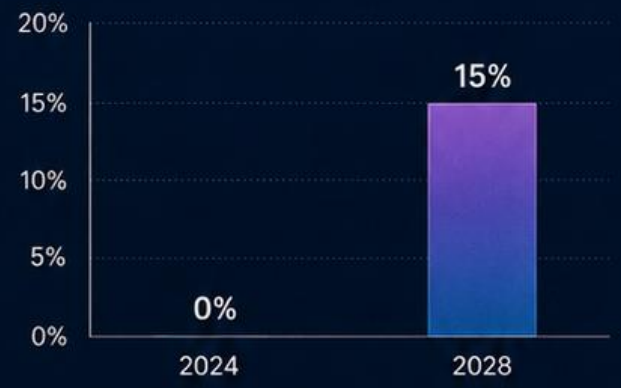

Day-to-day work decisions expected to be made autonomously by agentic AI.

Source: Gartner, June 2025.

Chart 4 · The Agentic Execution Gap

33%

Enterprise applications expected to include agentic AI by 2028

40%+

Agentic AI projects expected to be canceled by end of 2027

Source: Gartner, June 2025.

Why autonomy changes the control model

Identity and accountability

Know who and what is acting, on whose behalf, and how outcomes are attributed.

Connector and tool permissions

Control which systems and tools agents can access and what actions they can perform.

Data access and external action paths

Govern data use and manage every outbound action across internal and external systems.

Approval thresholds and human oversight

Define when to require human review, escalation, and exception handling.

Evidence, drift detection, and auditability

Capture complete evidence, monitor behavior, and detect drift early.

The goal is not simply to deploy more agents. It is to deploy **governed agents** that deliver value safely, consistently, and with clear accountability.

The Four AI Sprawl Types

AI risk is not one sprawl problem. Enterprises must manage usage, agents, identities, and tokens together.

1. AI Use Sprawl

Employees use external AI tools without centralized policy, review, or data handling controls.

2. Agent Sprawl

Business units deploy agents and automations that may lack inventory, ownership, purpose, or oversight.

3. NHI Sprawl

Non-human identities, service accounts, and machine credentials proliferate without clear accountability.

4. Token Sprawl

API keys, long-lived secrets, and application tokens create hidden access pathways for AI systems.

Chart 5 • Shadow AI Is a Material Security and Financial Issue

63%

Organizations lacking or still developing AI governance policy

20%

Breached organizations reporting a Shadow AI-related incident

$670K

Incremental average breach cost associated with Shadow AI

Source: IBM Cost of a Data Breach Report 2025, as reported by secondary analyses.

Shadow AI is not merely a policy violation.
It is an exposure, cost, and control-verification issue.

Executive White Paper · 2026

Why Policy and Inventory Alone Fail

The material AI risk often lives in the relationships between agents, identities, data, connectors, and external actions.

A list of agents is helpful but insufficient.
Enterprises need a graph of how AI systems are connected to models, identities, tools, data stores, business owners, and outbound actions.

The AI Data-to-Action Graph

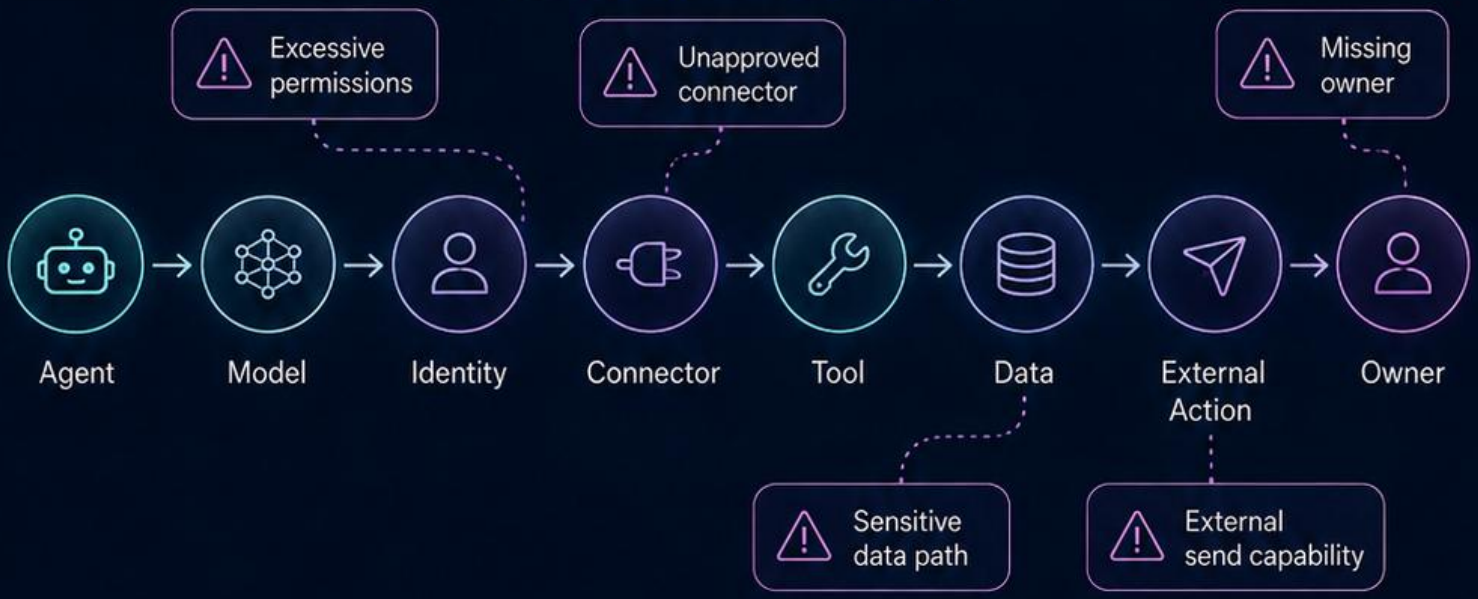

What governance must answer

1. What AI systems exist?
2. Who owns each one?
3. What identity does it run under?
4. What can it access and do?
5. What does it cost?
6. Can we prove policy and control?

Executive White Paper · 2026

Frameworks and Regulatory Pressure

A durable AI operating model must align to governance frameworks while producing evidence for emerging regulation.

Chart 6 · NIST AI RMF Mapping

NIST AI RMF	Govern360 Operating Model
GOVERN →	Ownership, purpose, policy, approval
MAP →	AI discovery, inventory, agents, identities, data paths
MEASURE →	AI Exposure Score, risk, cost, drift, evidence quality
MANAGE →	Policy compilation, enforcement, read-back, remediation workflow

Framework reference: NIST AI RMF 1.0 and Generative AI Profile. Product mapping: Govern360.

Chart 7 · EU AI Act Operating Timeline

1 Aug 2024
AI Act entered into force

2 Feb 2025
Prohibited AI practices and AI literacy obligations begin applying

2 Aug 2026
AI Act generally becomes applicable; Article 50 transparency obligations apply

Source: European Commission / EU Digital Strategy.

AI governance must become an evidence-producing operating model, not a future compliance project.

What a Real AI Estate Looks Like

The strongest governance signal often appears in the cross-agent join, not in a single agent viewed alone.

625 agents assessed

Live tenant example

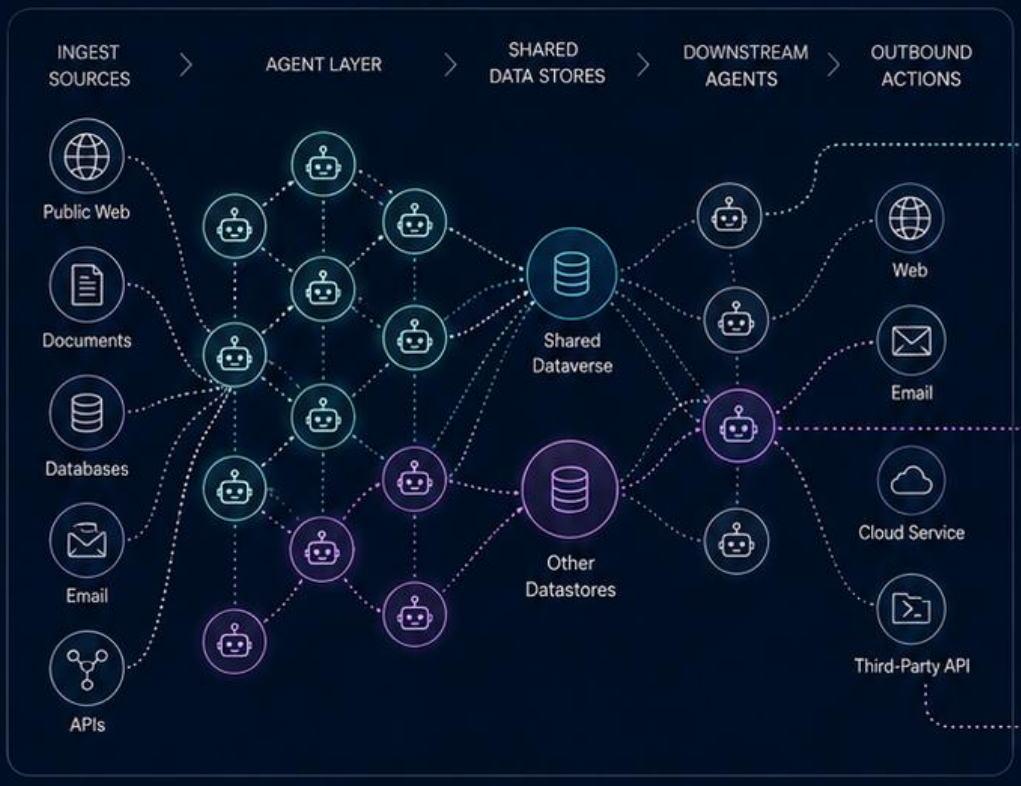

48 agents
Ingest unvetted content and write to shared Dataverse

1 downstream agent
Reads that shared datastore and can send externally

3 agents
Can send data outside the estate

Why this matters

The risk is not only in a single agent. It emerges when multiple agents, shared data stores, identities, and outbound actions create an exposure path across the estate.

These findings are from a live customer or internal tenant assessment in July–August 2026, anonymized and presented as an illustrative example. They are not representative of all enterprises and do not guarantee similar findings in every environment.

Executive White Paper · 2026

A 90-Day Adoption Roadmap

How enterprises can move from AI uncertainty to governed visibility and control.

1 2 3

PHASE 1

Days 0–30 · Discover

- Connect core SaaS, AI, identity, and data platforms
- Establish baseline inventory of agents, NHIs, and tokens
- Identify missing ownership and high-risk exposure paths

PHASE 2

Days 31–60 · Govern

- Define purpose, ownership, and policy requirements
- Prioritize high-risk agents and external actions
- Map spend, exposure score, and compliance evidence gaps

PHASE 3

Days 61–90 · Control and Prove

- Compile policy through existing control planes
- Establish remediation workflows and approval thresholds
- Produce evidence for security, finance, and compliance stakeholders

Methodology and publication notes

- External market statistics in this report are attributed to primary or official sources wherever possible.
- Gartner figures are cited with attribution and do not imply endorsement.
- IBM Shadow AI statistics are presented as reported by secondary analyses unless licensed source access is available.
- Customer observation data is anonymized and illustrative.

The next enterprise control challenge is not whether AI exists. It is whether every AI agent, identity, connector, and external action can be discovered, governed, and explained.