GOVERNANCE / RESEARCH BRIEF

Agent Governance Checklist

A board-ready operating checklist for making AI agents inspectable before they are scaled.

01	5 control domains
02	0 autonomous writes
03	1 owner per agent

CONTROL THESIS

Autonomy is rising faster than accountability.

Most agent initiatives are no longer blocked by model access. They are blocked by unclear ownership, weak permission boundaries, poor auditability, and unmeasured operating risk. The practical governance question is not whether an agent is intelligent. It is whether the business can prove who the agent is, what it was allowed to touch, why a recommendation was made, where a human approved it, and how performance is monitored over time.

01
Governance has to be designed into the workflow, not added after the demo.

02
Agent identity, risk class, permissions, and audit trail should survive process restarts.

03
The first useful artifact for a buyer is often a readable agent register, not a complex automation.

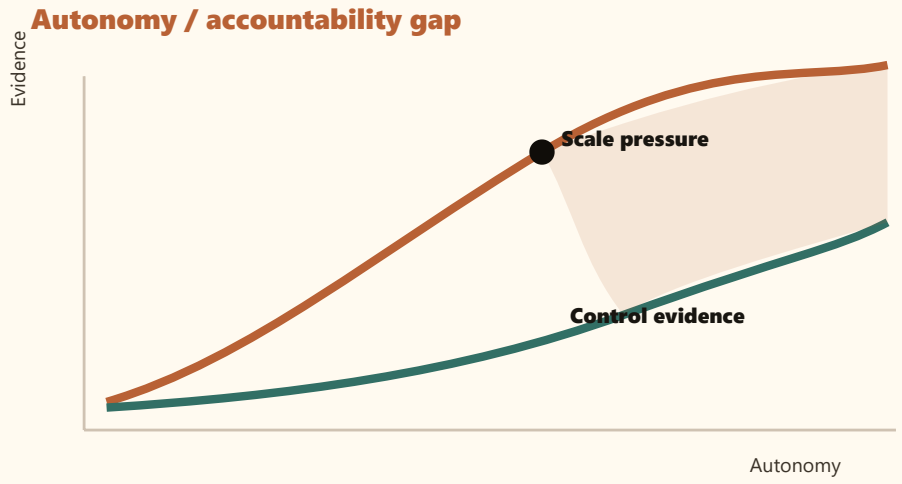

MCKINSEY & COMPANY / THE STATE OF AI: GLOBAL SURVEY 2025

AI use is broad, but scaling remains uneven. That gap makes governance and operating discipline a commercial differentiator.

MIT / CAMBRIDGE / STANFORD-AFFILIATED RESEARCH TEAM / THE 2025 AI AGENT INDEX

Agent transparency is becoming a research concern: buyers need to understand agent capabilities, dependencies, and safety boundaries.

MCKINSEY & COMPANY / STATE OF AI TRUST IN 2026: SHIFTING TO THE AGENTIC ERA

Trust maturity is improving, but agentic controls and governance remain lagging dimensions for many organizations.

AGENT CONTROL MATURITY HEATMAP

Governance controls the buyer can inspect.

Governance stack from identity to evidence: the buyer should be able to inspect every layer before autonomy increases. Read the artifact as a governance maturity map: identity comes first, then scoped authority, approval evidence, operating review, and only then higher autonomy.

Agent control maturity heatmap

IDENTITY
Assign a durable agent_id before first run and connect it to ownership, tool scope, memory namespaces, and audit entries.

AUTHORITY
Separate READ_ONLY, FINANCIAL, and DESTRUCTIVE actions so authority is explicit before the system can act.

APPROVAL
Hold consequential steps for human review and record approver, evidence, timestamp, and terminal status.

EVIDENCE
Preserve audit records and CLEAR scorecard outputs so the system can be reviewed by operators, risk owners, and legal counsel.

GOVERNANCE WORKSHOP

Turn agent sprawl into an inspectable control register.

Use the meeting to locate invisible agents, assign owners, classify risk, and decide which actions stay blocked until evidence exists.

INSPECTION QUESTIONS

- Which person owns this agent, its purpose, and its revocation decision?
- What tool calls can the agent make without approval, and which are blocked?
- Where does the audit trail prove prompt, data source, tool call, and approval status?
- Which metric would trigger rollback, review, or reduced autonomy?

DO NOT SHIP UNTIL

- Do not treat a model policy as a permission system.
- Do not let teams deploy agents without a revocation path.
- Do not sell compliance certification; provide evidence for qualified review.

DAY 1	Create the agent register and classify every known agent, copilot, and automation.	Agent identity table, owner, purpose, risk class, and revocation field.
WEEK 1	Map tool permissions and human approval gates for the first commerce workflow.	Tool/action matrix with READ_ONLY, FINANCIAL, and DESTRUCTIVE boundaries.
WEEK 2	Run the first CLEAR review using synthetic or historical workflow examples.	Cost, latency, efficiency, accuracy, and reliability baseline.

SOURCE LEDGER

Source ledger and governance boundary

The brief combines external research signals with HARNEXA's implementation boundary: identity, permission, approval, audit, CLEAR evaluation, and no autonomous public execution.

MCKINSEY & COMPANY

The State of AI: Global Survey 2025

McKinsey reports broad AI adoption, but a much smaller share of organizations scaling AI programs and agentic AI.

https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

MIT / CAMBRIDGE / STANFORD-AFFILIATED RESEARCH TEAM

The 2025 AI Agent Index

The index reviews prominent deployed AI agents and highlights the need to expose origins, capabilities, ecosystem dependencies, and safety features.

https://aiagentindex.mit.edu/

MCKINSEY & COMPANY

State of AI trust in 2026: Shifting to the agentic era

McKinsey reports that responsible AI maturity is improving, but strategy, governance, and agentic AI controls still lag.

https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/state-of-ai-trust-in-2026-shifting-to-the-agentic-era

RESEARCH STANCE

External research is cited as market signal; HARNEXA interpretation is kept separate from source claims.

BUYER ARTIFACT

Every report includes a concrete visual artifact that can be inspected in a sales, risk, or architecture review.

OPERATING BOUNDARY

Persistent identity, owner, permissions, approval, audit, CLEAR. The asset is for governed review, not autonomous production execution.

HARNEXA BOUNDARY

Technical governance aid only. Not legal advice or compliance certification. HARNEXA AI builds audit-ready foundations for qualified review. No PDF in this library claims EU AI Act certification, legal advice, payment readiness, or autonomous production execution.