HIVE
ADVISORY GROUP

A HIVE ADVISORY GROUP BRIEF

Ten steps to *AI governance*

What the leading strategy and advisory firms agree on, consolidated for mid-market leadership teams.

October 2026

hiveadvisorygroup.com

WHY THIS BRIEF

Eight firms, one broadly shared sequence

Over the past two years McKinsey, BCG, Bain, Deloitte, PwC, EY, KPMG and Gartner have each published research and guidance on governing AI. The reports differ in emphasis and vocabulary, but they agree on more than they disagree. Read together, they describe much the same sequence of work.

This brief consolidates that guidance into ten steps. For each one it adds a view on what the step looks like in a mid-market company, where there is rarely a dedicated AI risk team and the same few leaders carry the work alongside their existing roles.

1 in 3

About a third of organizations reach level three or higher, on a four-level scale, in AI strategy, governance and controls on AI agents.

McKinsey, AI Trust Maturity Survey, 2026[10]

21%

Of companies planning to deploy AI agents within two years, the share that report a mature model for governing them.

Deloitte, State of AI in the Enterprise, 2026[5]

40%

Of employees say their workplace has a policy or guidance on the use of generative AI.

University of Melbourne and KPMG, 2025[8]

A note on the data: most of these surveys draw mainly on large enterprises, and the mid-market is measured less often. Each figure in this brief is attributed to its source, with the survey population noted where it matters.

AT A GLANCE

1	**Ownership at the top**	Who answers for AI in this business?
2	**A stated risk appetite**	What will we accept, and what is off the table?
3	**An inventory of the AI already in use**	Where is AI already at work here?
4	**A reconciliation against commitments**	What have we already promised in writing?
5	**Risk tiers for use cases**	Which uses need the most attention?
6	**A policy people can follow**	Do employees know what is expected of them?
7	**Terms for vendors and third-party AI**	What are suppliers doing with AI, and with our data?
8	**Controls for AI agents**	What may an agent do, and who signed off?
9	**Monitoring, testing and an incident plan**	How would we know if something went wrong?
10	**A review rhythm, with tooling chosen to fit**	How does this keep pace as AI changes?

1 Ownership at the top

Bain is direct on this point: AI risk belongs to the CEO, who sets the appetite for it and is answerable for the results.[1] McKinsey's 2026 AI Trust Maturity Survey of about 500 organizations found that those with a clearly accountable owner for responsible AI averaged 2.6 on a four-level maturity scale, against 1.8 for those without one.[10] PwC recommends a three lines model, so that the people who build AI, the people who review it and the people who assure it each have a defined role.[13]

> **AT MID-MARKET SCALE**
>
> A new department is rarely needed. What tends to work is one named executive who is accountable, supported by a small group drawn from operations, technology, legal and finance that meets on a set schedule.

2 A stated risk appetite

Bain describes risk appetite for AI agents as an ongoing discussion between the CEO and the board rather than a one-time approval, tied to how mature the organization's controls are today.[2] BCG groups the risks into three families that give that conversation some structure: quality (inaccurate or low-quality output), safety (bias and harmful content) and security (data leakage, manipulation and fraud).[3]

> **AT MID-MARKET SCALE**
>
> This can be a single page. It usually covers the decisions AI may support but not make, the categories of data that stay out of AI tools, and the level of error leadership is prepared to accept in lower-stakes work.

3 An inventory of the AI already in use

Every firm starts in the same place. Bain's summary is that a company cannot govern an AI estate it has not inventoried.[1] BCG suggests a 30-day inventory of every agent, including pilots, vendor tools and deployments led by business teams.[4] McKinsey describes a centrally managed portfolio that records the owner, the use case, the data each system can reach and its status.[9] Visibility is a common gap: in EY's survey of 975 C-suite leaders at large companies, half said they do not have a high level of visibility into how employees use AI agents.[6]

> **AT MID-MARKET SCALE**
>
> In our experience the inventory is longer than leadership expects. Much of it is AI that arrived inside software the company already pays for, through vendor updates nobody was asked to approve, alongside the tools employees have adopted on their own.

4 A reconciliation against commitments

The published guidance covers the regulatory side of this step. McKinsey advises identifying the rules that apply to each use and preparing for likely requirements such as human oversight, data protection and fairness.[9] BCG points to the EU AI Act's four risk levels as a reference for classifying systems.[3]

> **HIVE'S ADDITION**
>
> We would add a second reconciliation that receives less attention: what the business has already promised in writing. Client contracts, engagement letters and supplier agreements can carry obligations about how work is done and how information is handled. Setting the inventory against those documents is where we most often find exposure.

5 Risk tiers for use cases

Not every use of AI warrants the same scrutiny. McKinsey's Rich Isenberg describes tiering by autonomy: a knowledge assistant mainly raises accuracy questions, an agent that approves invoices needs more human involvement, and a fully autonomous agent needs controls that keep it within its designed boundaries.[11] BCG suggests ranking agents by the likelihood and impact of exceeding their authorization, with priority for those that touch customer data, regulated information, external communications, money movement or production systems.[4] Deloitte reports that the companies seeing the most success with agents begin with lower-risk uses and expand deliberately.[5]

> **AT MID-MARKET SCALE**
>
> Three tiers are usually enough: uses that can proceed under the general policy, uses that need a named owner and a review, and uses that need leadership sign-off.

6 A policy people can follow

In the University of Melbourne and KPMG study of more than 48,000 people across 47 countries, 66% of employees said they rely on AI output without evaluating its accuracy and 56% said they had made mistakes in their work because of AI.[8] McKinsey's trust survey found that nearly 60% of organizations cite knowledge and training gaps as the main barrier to responsible AI, up from about 50% a year earlier.[10] Isenberg adds that where guardrails are optional, shadow agents and inconsistent standards tend to follow.[11]

> **AT MID-MARKET SCALE**
>
> Short policies get read. Two pages covering the approved tools, the data that stays out of them, when a person checks the output, and a quick route for requesting something new tend to do more than a long document. Training carries as much weight as the text.

7 Terms for vendors and third-party AI

McKinsey advises updating third-party risk management to cover agents, and holding outside agents to the same security, governance and ethical standards as internal systems.[9] Bain makes a related point: risk has to be assessed across the whole system, including the models, tools and connectors a company does not own, because a weakness anywhere in that supply chain passes to the company.[2]

> **AT MID-MARKET SCALE**
>
> This is often the largest part of the picture, because much of a mid-market company's AI is bought rather than built. The questions that tend to matter in supplier agreements are what the vendor does with company data, whether that data trains the vendor's models, how new AI features are introduced, and what the company is told when something goes wrong.

8 Controls for AI agents

Agents act, which changes the control problem. Deloitte found that close to three-quarters of companies plan to deploy agentic AI within two years, while only 21% of those companies report a mature model for governing agents.[5] In EY's survey, two-thirds of large companies allow employees to build or deploy agents in some form, and 60% of those have formal, organization-wide policies for it.[6] The firms converge on the controls: a registry of every agent, a distinct identity and minimum permissions for each, a named human owner, and human confirmation before high-stakes or irreversible actions.[2, 4, 11]

> **AT MID-MARKET SCALE**
>
> A company with a handful of agents can keep this simple: a list of the agents, who owns each one, what each can reach, and which actions wait for a person.

9 Monitoring, testing and an incident plan

In PwC's survey, half of executives named operationalizing responsible AI as their biggest hurdle. PwC's recommendation is to automate testing and monitoring across the AI lifecycle.[13] EY reports that large companies with real-time monitoring were 34% more likely to report improved revenue growth and 65% more likely to report improved cost savings.[6] McKinsey and Bain both recommend a record of what an AI system did and why, a way to stop it, and a fallback worked out in advance. Bain adds that a stop mechanism only counts once it has been tested.[2, 9]

> **AT MID-MARKET SCALE**
>
> The plan can be modest: periodic spot checks on the highest-tier uses, a log of issues, one person to call when something looks wrong, and an agreed way to pause a tool.

10 A review rhythm, with tooling chosen to fit

PwC describes responsible AI as something to be reassessed continually as technology and risks change.[13] The reference frameworks move too: NIST is revising its AI Risk Management Framework, first released in January 2023.[12] On tooling, two findings are worth holding together. In a Gartner survey of 360 organizations, those using AI governance platforms were 3.4 times more likely to achieve high governance effectiveness.[7] The authors of McKinsey's trust survey add that these capabilities cannot be acquired as a product, and that they depend on policies, processes and people as well as technology.[10]

> AT MID-MARKET SCALE
>
> A quarterly review of the inventory, the tiers and the incident log is a reasonable rhythm. For companies comparing governance tools, the earlier steps tend to make the choice simpler, because they define what the tool needs to do.

SEQUENCING

One way to order the first 90 days

BCG frames agent governance as a 30, 60 and 90-day sequence.[4] The same shape suits the ten steps, with the tenth becoming a standing rhythm once the first nine are in place.

DAYS 1 TO 30

Steps 1 to 3

An owner is named, the risk appetite is written down and the inventory is complete.

DAYS 31 TO 60

Steps 4 to 6

Usage is reconciled against commitments, uses are tiered and the policy is issued.

DAYS 61 TO 90

Steps 7 to 9

Vendor terms are reviewed, agent controls are in place and monitoring has started.

FOR TEAMS USING A FORMAL FRAMEWORK

How the steps line up with NIST

The NIST AI Risk Management Framework is voluntary and organizes the work into four functions.[12] The mapping below is our own reading and is approximate.

NIST FUNCTION	STEPS IN THIS BRIEF
Govern	1 Ownership, 2 Risk appetite, 6 Policy, 7 Vendor terms
Map	3 Inventory, 4 Reconciliation, 5 Risk tiers
Measure	9 Monitoring and testing
Manage	8 Agent controls, 9 Incident plan, 10 Review rhythm

About Hive Advisory Group

Hive is an AI strategy and governance advisory firm based in Nashville, working with mid-market companies across Tennessee. Each engagement is led by a senior team assembled for the work, with one named lead. Our AI Risk & Governance Review runs four weeks and covers steps 3 to 5 of this brief, with a board-ready roadmap for the rest.

TALK IT THROUGH

A 15-minute call is usually enough to know whether there is something worth pursuing.

MORE ON THE REVIEW

hiveadvisorygroup.com/ai-risk-governance-review

SOURCES

1 **Bain & Company.** "What Is AI Governance?" July 2026. bain.com

2 **Bain & Company.** "Agentic AI Governance, Risk, and Controls for Business Leaders." July 2026. bain.com

3 **Boston Consulting Group.** "Managing Risks to Accelerate the AI Transformation." December 2024. bcg.com

4 **Boston Consulting Group.** "The Authorization Gap: Why Yesterday's Controls Won't Work with Today's Agents." September 2026. bcg.com

5 **Deloitte.** "The State of AI in the Enterprise: The Untapped Edge." January 2026. Survey of 3,235 business and IT leaders in 24 countries. deloitte.com

6 **EY.** Responsible AI Pulse survey, second phase. October 2025. Survey of 975 C-suite leaders in 21 countries at organizations with more than US$1 billion in annual revenue. ey.com

7 **Gartner.** "Global AI Regulations Fuel Billion-Dollar Market for AI Governance Platforms." Press release, February 17, 2026. gartner.com

8 **University of Melbourne and KPMG.** "Trust, attitudes and use of artificial intelligence: A global study 2025." April 2025. Survey of more than 48,000 people in 47 countries. kpmg.com

9 **McKinsey & Company.** "Deploying agentic AI with safety and security: A playbook for technology leaders." October 2025. mckinsey.com

10 **McKinsey & Company.** "State of AI trust in 2026: Shifting to the agentic era." March 2026. Survey of about 500 organizations. mckinsey.com

11 **McKinsey & Company.** "Trust in the age of agents." The McKinsey Podcast, with Rich Isenberg. March 2026. mckinsey.com

12 **National Institute of Standards and Technology.** AI Risk Management Framework (AI RMF 1.0). January 2023. nist.gov

13 **PwC.** "2025 Responsible AI survey: From policy to practice." October 2025. Survey of 310 US business leaders. pwc.com

This brief summarizes and interprets published research by the organizations named above. Hive Advisory Group is independent and is not affiliated with, sponsored by or endorsed by any of them. Figures are as reported by each source at the time of publication, and survey populations differ. The brief is general information, not legal advice. © 2026 Hive Advisory Group.