



Beyond the pilot: Building, scaling and sustaining agentic AI initiatives

Table of contents

An ongoing capability rather than a one-time initiative	03
The balancing act: Navigating strategic dilemmas	04
1. Funding model: Enabling innovation at scale	05
2. Technology ownership: Creating autonomy within the governed scale	05
3. Solution strategy: Build custom versus buy Commercial-off-the-Shelf (COTS)	05
4. Talent strategy: Evolving digital workforce	06
5. AI ethics and governance: Imbibing transparency and bias checks	06
6. Business orchestration: Balancing internal control and external leverage	06
Pivoting from strategy to execution: Scaling the right way	07
Sustaining agentic AI for long-term value	10
Closing thoughts	11



An ongoing capability rather than a one-time initiative

Agentic AI is transforming the way work is done. It helps teams make faster decisions, automate routine tasks and improve outcomes across functions.

In India, many organisations have started using AI agents. Pilots are active across sectors, and real value is achieved when these agents are built, scaled and sustained as part of the business design. Without a clear enterprise-wide strategy, AI often remains fragmented. This leads to higher costs, fragile systems and governance challenges.

About 79 percent¹ of organisations say they use AI agents in some form. Only 19 percent¹ have deployed them at scale. Just 5 percent² of integrated pilots are delivering multi-million-dollar value. The models are working as expected. What is needed now is more precise coordination on funding, ownership, talent and governance.

CXOs face six critical dilemmas on the path from pilots to scalable platforms:



For Indian enterprises, this is a moment to lead. With the right architecture, operating model and talent strategy, agentic AI can unlock speed, scale and smarter decisions. It requires strategy, tools and experience to build what lasts.

The next step is to explore how agents can become part of your business, contributing alongside your teams and helping shape outcomes that matter.

¹ Multimodal.dev

² MIT Nanda report – State of AI in Business, July 2025

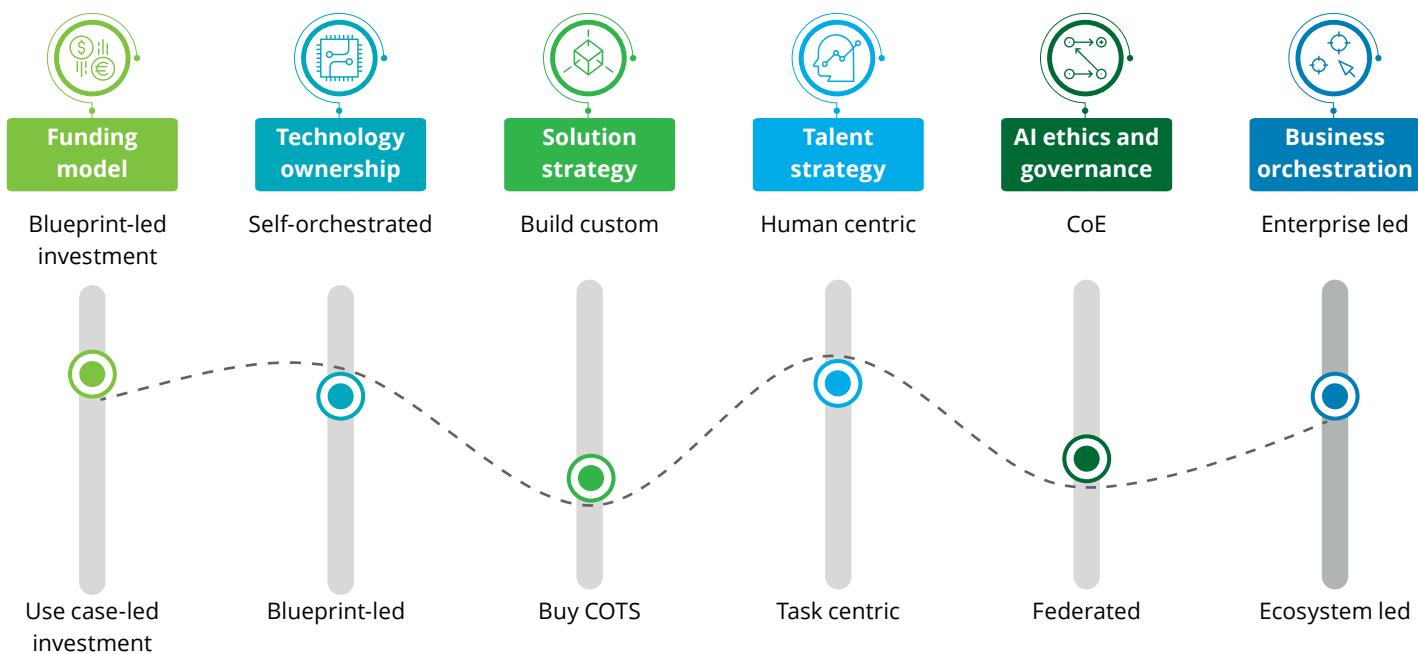


The balancing act: Navigating strategic dilemmas

Nearly seven in ten leaders³ report that their AI projects never reach full operations. Agentic AI is no exception. Scaling agentic AI is a multifaceted journey. It involves engaging with six strategic levers, each presenting a unique dilemma that benefits from deliberate and sustained effort rather than quick

solutions. These dilemmas require foresight and deliberate choices (Figure 1). When these levers are balanced, agentic AI moves beyond pilots and Proofs Of Concept (PoC) to become a source of competitive advantage.

Figure 1: The balancing act: Navigating six key trade-offs in agentic AI adoption



³ Blue Prism Global AI Survey, 2025

Treating these levers as isolated decisions can lead to fragmented budgets, redundant tools, compliance gaps and adoption fatigue.

To overcome these pitfalls, our approach emphasizes a coordinated strategy across six levers, anchored in

the enterprise context and powered by global delivery capabilities.

This alignment ensures that agentic AI initiatives are technically strong and strategically well-integrated, with the potential to scale effectively and remain sustainable over time.

1

Funding model: Enabling innovation at scale

The dilemma: A blueprint platform-led investment model delivers economies of scale through shared assets such as data lakes and model registries, yet it often delays time-to-value. A use case-led approach demonstrates ROI quickly, although it risks creating silos of tools and data.

What should you do: Adopt a two track funding approach. Reserve most of the AI budget for a central platform that establishes data, compute and governance foundations. Allocate the rest to agile, ROI gated pilots that feed demand into the platform.

Dual track AI adoption (combining experimentation and platform investment) can deliver up to 3.2x faster model scaling and 27 percent lower Total Cost of Ownership (TCO) compared with only piloting or only scaling approaches.⁴

2

Technology ownership: Creating autonomy within a governed scale

The dilemma: Self-orchestrated but unchecked autonomy accelerates experimentation but may compromise compliance, security and scalability. Strict blueprint-led centralisation reduces risk but can slow delivery.

What should you do: Implement guardrails first governance as a minimal, non negotiable set of standards (e.g., security, data residency, privacy, drift monitoring). Teams may then self orchestrate everything else.

3

Solution strategy: Build customised solutions versus buy Commercial-off-the-Shelf (COTS)

The dilemma: In-house, custom-built solutions (models, orchestration) protect strategic IP but are expensive and consume engineering bandwidth. COTS platforms and agents accelerate time to market; however, they lock-in vendors and limit differentiation.

What should you do: Make balanced investment choices. Buy for parity, build for differentiation. Adopt market tools for generic functions (retrieval, external research, summarisation). Develop custom agents with domain logic, proprietary data connectors, policies and workflows that encode how your business operates.

⁴ Forrester AI Spending Tracker, 2025



Talent strategy: Evolving digital workforce

The dilemma: Human centric models boost employee engagement and trust but require upskilling. Task-centric models drive cost efficiency, though they can undermine morale and raise compliance concerns.

4

What should you do: Develop a culture of an AI-mindset where employees see agents as collaborators. Focus on enhancing human capabilities with AI, while automating repetitive, routine tasks for greater efficiency. Establish Human-in-the-Loop (HITL) checkpoints for high-risk decisions to maintain oversight, trust and accountability.

About 69 percent⁵ of leaders recognise the importance of reinventing the Employee Value Proposition (EVP) to reflect increased human-machine collaboration.

5

AI ethics and governance: Imbibing transparency and bias checks

The dilemma: A single Center of Excellence (CoE) ensures consistency, but it can also become a bottleneck. Federated governance models offer speed, yet they can lead to inconsistent safeguards.

What should you do: Adopt a hybrid governance model. The CoE defines baseline ethical controls (e.g., explainability standards, audit log requirements) that each function must adopt, while allowing units to add context specific extensions (e.g., fairness definitions for credit scoring for lending underwriting workflows).

6

Business orchestration: Balancing internal control and external capacity

The dilemma: Complete internal control of processes (workflows, data) safeguards IP but slows innovation. Using external ecosystems (agentic AI platforms and pre-built agents) accelerates capability acquisition, although it raises concerns around data sovereignty and compliance.

What should you do: Keep critical processes, such as those tied to revenue or regulatory risk, within the company, while using trusted external agents for low-risk, support tasks. Use a simple scoring model to determine which workflows should remain internal and which can be safely outsourced. This helps protect sensitive data while still speeding up innovation through external collaborations.

⁵ Deloitte Research



Pivoting from strategy to execution: Scaling the right way

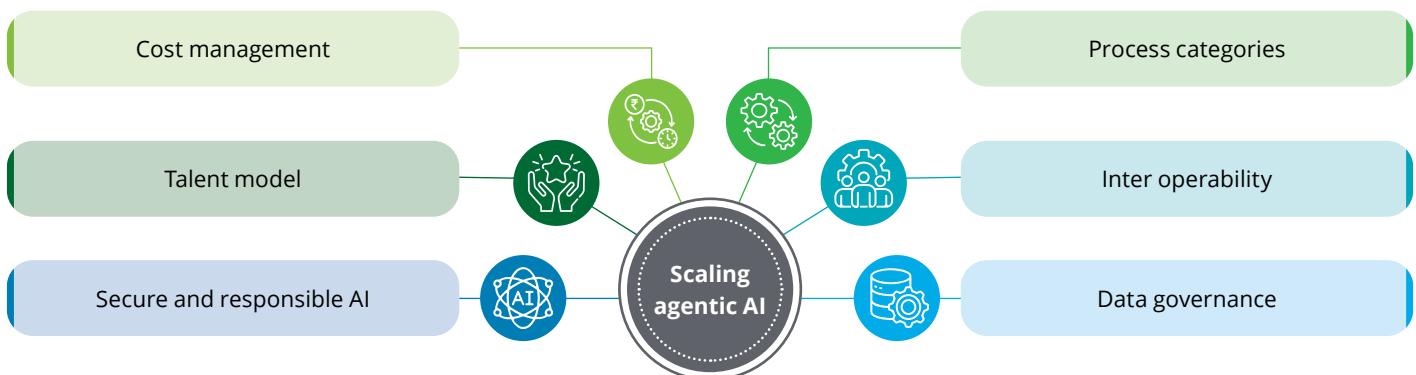
With the six levers and their decisions defined, the next challenge is turning strategy into execution. Many organisations stall here; pilots multiply but rarely scale, and execution is fragmented across teams.

Organisations need an operating model that unites architecture, data, controls and people into one way of working. For sharp, enterprise-grade execution, these six

guiding principles provide a foundation for scale, ensuring that agentic AI moves beyond demos to deliver measurable business outcomes (Figure 2).

Agentic AI becomes transformative when treated as a horizontal capability supported by governance, common infrastructure and trusted expertise.

Figure 2: Six key guiding principles to scaling agentic



1

Target process categories

Agent success starts by addressing a class of problems rather than isolated use cases. The focus should be on workflows where agents solve process-driven business challenges. Many organisations stall by perfecting individual solutions. Instead, validate quickly through Proof of Technology → PoC → Scale, letting processes and use cases guide technology choices rather than chasing the latest AI features. This approach accelerates adoption and ensures impact at scale.

Ask yourself: Which class of problems matter most, and how will we measure improvement?

What should you do

- Start small, think big. Deliver the smallest solution that proves value.
- Focus on classes of problems (such as customer support or financial reconciliation) rather than one-off use cases.
- Define guiding principles early, identify generic agents and reusable components, and set up a shared repository to avoid duplicative agent builds across the organisation.

2

Build an interoperable agent network

Create agent ecosystems that work together seamlessly across your entire organisation. Multi-agent workflows fail without consistent contracts. Agents must hand off tasks, share context and call tools without rework.

Ask yourself: How do we prevent new silos while dismantling old ones?

What should you do

- Prioritise interoperability for core business processes, allow customisation for specialised functions that do not need integration.
- Establish shared authentication frameworks for cross-departmental agent communication.
- Implement standardised data exchange protocols across systems and vendor platforms.

3

Govern data for accessibility and autonomy

Agents are only as good as what they can retrieve. Make data and documents reliable, current and easy for agents to find. Think shared data lake with specialised agents for sales, operations and finance, each accessing common data but optimised for their specific tasks.

Ask yourself: How do we enforce data control while allowing autonomy?

What should you do

- Build a central data and compute foundation with task-specific agents.
- Build modular architectures that let you switch vendors as technology evolves, even if initial setup takes longer.

4

Embed secure and responsible AI by design

Autonomy requires guardrails and a core operating framework. Security, privacy and risk controls must be embedded into design, release and monitoring processes from the start. Observability is critical to ensure that agent behaviour, performance and decision-making can be traced, audited and understood in real time. AgentOps then provides the production capability to run agents safely, delivering continuous monitoring, policy enforcement, automated testing and rollback mechanisms that ensure controlled autonomy at scale.

Ask yourself: What level of autonomy is acceptable, and what checks are required before scale?

What should you do

- Implement a two-tier governance structure:
 - Central Executive Council (CEC) handles strategy and major risk decisions.
 - Design authority directs consistent standards, patterns, libraries and tooling choices.
- Define a framework for autonomy levels with the required approvals and human-in-the-loop points for each level. Make these conditions mandatory for release.
- Integrate observability tools, such as real-time dashboards, log aggregation and performance metrics, to monitor agent health, detect drift and ensure compliance with operational and ethical standards.
- Build AgentOps-as-a-Capability for tracing, alerts, incident playbooks and controlled rollback so teams can scale with confidence.

5

Transform talent model

Capacity comes from a small central core and business-led squads that build and run agents close to the work. Solve talent gaps through flexible workforce planning by borrowing specialised expertise on demand, moving talent across teams and hiring selectively for critical long-term roles. The planning must be managed through centralised budgeting that enables agile allocation based on project needs. Our studies confirm that over 75 percent⁶ of leaders believe skill-first decisions reduce bias and open opportunities to more people.

Ask yourself: Which skills remain central, and which ones do business teams use to move faster?

What should you do

- Create a central team that owns patterns, tools, training and approval gates. Hold them accountable for reuse and quality, rather than the volume of builds.
- Form cross-functional squads in business units to deliver use cases. Start with upskilling the existing workforce and fill only the gaps that truly need hiring.
- Use outsourced talent for specialised skills.

6

Track task-level AI economics

AI spend is variable and should track value. Make cost visible, predictable and steerable at the task level.

Ask yourself: What is the unit cost per task, and how will we stay within budget while improving outcomes?

What should you do

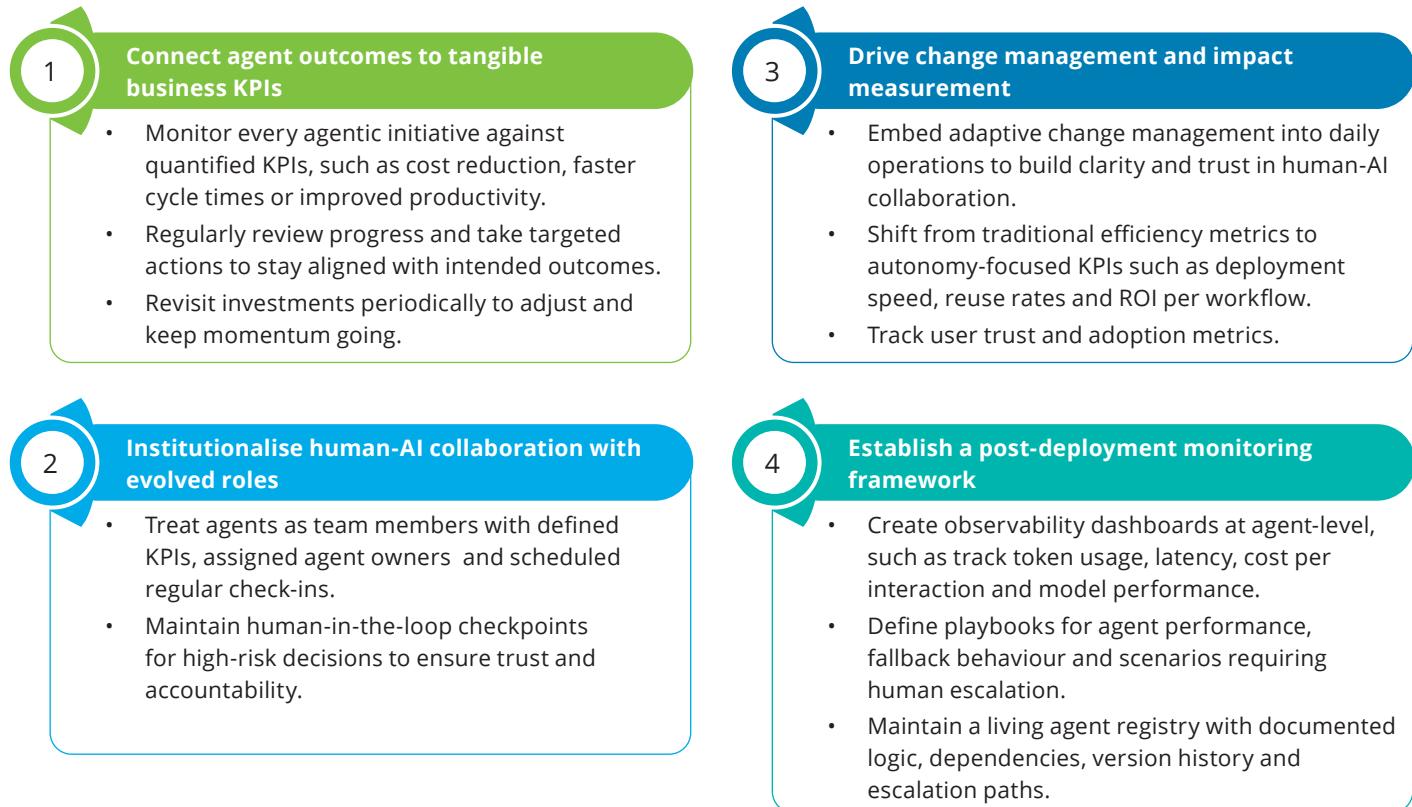
- Track spend by agent, team and task. Set budgets and limits, and review cost against value delivered every month.
- Monitor compute and LLM costs closely and demand transparent pricing to avoid unexpected overruns as usage grows.
- Define a modular structure to change models or platforms without business disruptions.



Sustaining agentic AI for long-term value

Agentic AI should be seen as a dynamic capability that evolves with changing business goals, technology trends and workforce needs. Sustaining this capability requires embedding adaptability, governance and measurement into the enterprise fabric.

Here is what leaders should do:



5

Implement continuous evaluation and drift detection

- Use automated pipelines to test agents against real-world benchmarks.
- Monitor for prompt drift, model degradation or context misalignment.
- Refresh prompts, workflows and domain logic based on business events and data updates to prevent degradation and bias.

6

Set up a responsible AI oversight loop

- Form a cross-functional review board to audit agent decisions, especially in sensitive domains.
- Include ethics, compliance and business stakeholders in quarterly reviews.
- Create feedback capture mechanisms and route feedback into retraining and prompt engineering workflows.

Closing thoughts

Agentic AI goes beyond automation, creating a clear pathway to lasting competitive advantage. When built with strategic clarity, scaled operational discipline and sustainable changes in ways of working, agents become value drivers across business functions. Leaders should move past one-off AI experiments, dissolve silos and build agent networks that work together across teams and systems to unlock real business value.

Winning in the age of agentic AI comes down to one critical question, and the rest follows:

“

Are you ready to move beyond pilots and build an enterprise where agents drive real value?

”

Connect with us

Ashvin Vellody

Partner, Deloitte India
ashvinv@deloitte.com

Moumita Sarker

Partner, Deloitte India
msarker@deloitte.com

Contributors

Akshay Heda

Anup Premarajan

Acknowledgements

Mohammed Riyaz Ahmed

Neha Kumari

Ruchira Thakur

Sagarika Gupta

Srishti Deoras

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte Asia Pacific Limited is a company limited by guarantee and a member firm of DTTL. Members of Deloitte Asia Pacific Limited and their related entities, each of which is a separate and independent legal entity, provide services from more than 100 cities across the region, including Auckland, Bangkok, Beijing, Bengaluru, Hanoi, Hong Kong, Jakarta, Kuala Lumpur, Manila, Melbourne, Mumbai, New Delhi, Osaka, Seoul, Shanghai, Singapore, Sydney, Taipei and Tokyo.

This communication contains general information only, and none of DTTL, its global network of member firms or their related entities is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication.

© 2025 Deloitte Touche Tohmatsu India LLP. Member of Deloitte Touche Tohmatsu Limited