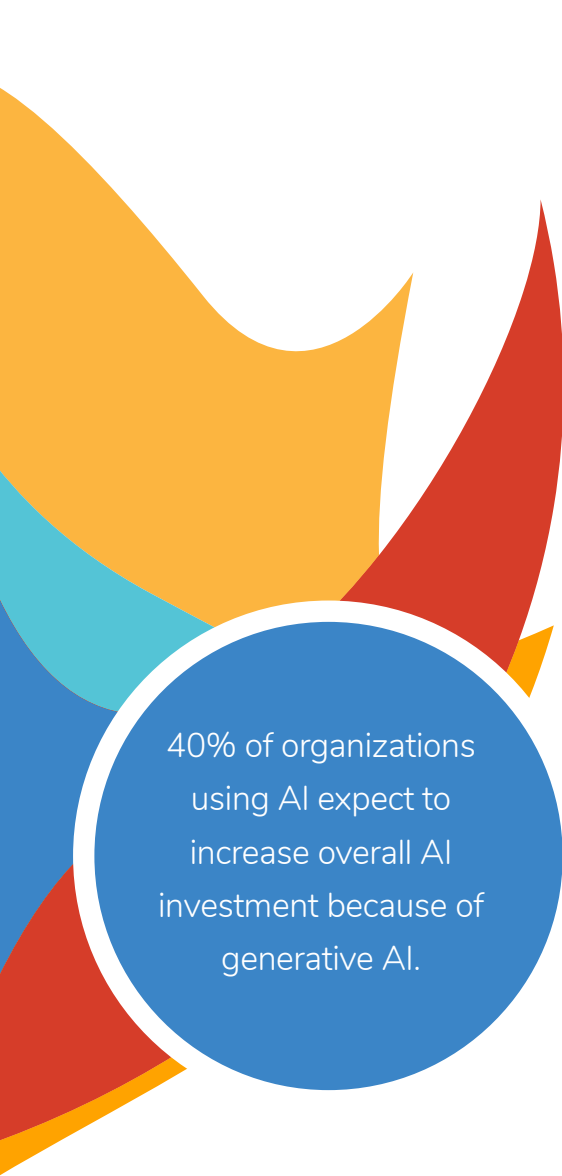




The Enterprise AI Shift

FROM EARLY EXPERIMENTS TO THE NEXT FIVE YEARS OF GOVERNED, ORCHESTRATED INTELLIGENCE



40% of organizations using AI expect to increase overall AI investment because of generative AI.

The release of ChatGPT in late 2022 triggered one of the fastest technology adoption cycles in enterprise history. Many organizations launched pilots or prototypes. According to [McKinsey's](#) “The state of AI in 2023: Generative AI’s breakout year,” in less than one year after many generative AI tools debuted, one-third of survey respondents said their organizations were already regularly using generative AI in at least one business function. Additionally, 40 percent of respondents whose organizations used AI said they expected to increase overall AI investment because of generative AI.

For many organizations, the first deployments were low-risk, highly visible prototypes, such as customer support chatbots designed to deflect basic inquiries, answer FAQs, and reduce call-center volume. From there, early AI experimentation rapidly expanded into other functional areas.

Since then, AI use has steadily evolved, and a noticeable shift has emerged in how businesses think about adopting it in a safe, manageable, and scalable way. This eBook examines what early AI use has taught us and how those insights are shaping the future of enterprise adoption.



What Organizations Learned from the First Wave


The first broad deployment of generative AI showed that while there were some employee productivity improvements, models alone were not enough. Early pilots revealed that when AI ran in isolation—disconnected from systems of record, business rules, and workflows—results were inconsistent and hard to trust. Teams learned that real value only emerged when AI was embedded into end-to-end processes, governed by clear logic, and aligned with how the business actually operates.

Organizations also discovered that AI amplified the strengths and weaknesses of their knowledge and data. Where documentation was current, structured, and curated, AI could surface generally good answers quickly. Where knowledge was fragmented or outdated, AI confidently produced the wrong guidance. This is in part why the first wave of AI was limited to experimentation and was not implemented into high value and mission critical processes. This also pushed leaders toward more deliberate knowledge management and highlighted the need for an orchestration layer that can route, select, and ground models in the right data at the right moment.

The first wave made it clear that human judgment still matters. When AI was allowed to make decisions without expert oversight, it sometimes missed nuances, such as policy exceptions, regulatory subtleties, or customer context, that experienced employees handle every day. The best outcomes came from designs where [AI accelerates and orchestrates](#) when to involve humans for review, escalation, or exception handling, rather than attempting to fully replace them.

Enterprises also ran into the limits of non-deterministic systems, those that cannot guarantee the same outcome despite having the same inputs. The same AI prompt could yield different answers within





AI's second wave isn't about more models—it's about orchestration, governance, and control.

the same LLM model, which is unacceptable in regulated or high-stakes decisions. To restore predictability and explainability, organizations began pairing probabilistic models with deterministic rules engines that enforce policy, eligibility, and compliance. In effect, the rules engine became the backbone that ensures AI-driven decisions are repeatable, auditable, and aligned with business logic.

Finally, early experimentation with agents and point solutions with embedded AI underscored the risks of uncoordinated AI. Separate pilots in different departments led to overlapping tools, inconsistent standards, and surprises in production behavior. This experience made clear that enterprises need a unifying orchestration layer: one place to define guardrails, connect systems, enforce rules, manage agents, monitor performance, and coordinate human and AI work across the organization.

The Second Wave of Enterprise AI

Many early AI efforts delivered pockets of value but failed to translate into organization-wide transformation. A second wave is now emerging in which enterprises are shifting from fragmented pilots toward a unified model built on governance, orchestration, and architectural integration. [Gartner](#) refers to the platforms for this stage as business orchestration and automation technologies (BOAT).

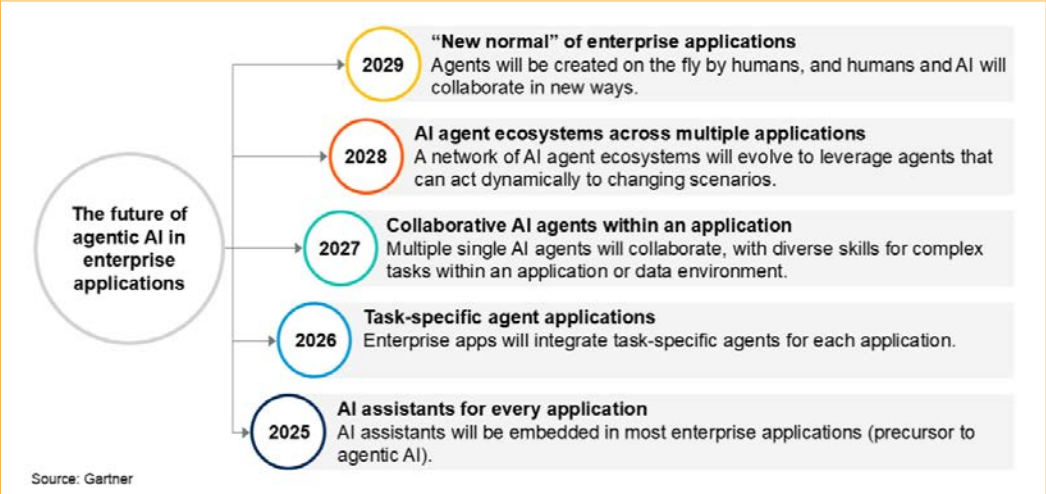
This phase is defined less by experimentation and more by operational discipline. Organizations are realizing that to use AI reliably, it must be connected to trusted systems, shaped by robust business rules, and embedded within workflows that drive consistent outcomes. At the same time, the rapid growth of models, copilots,

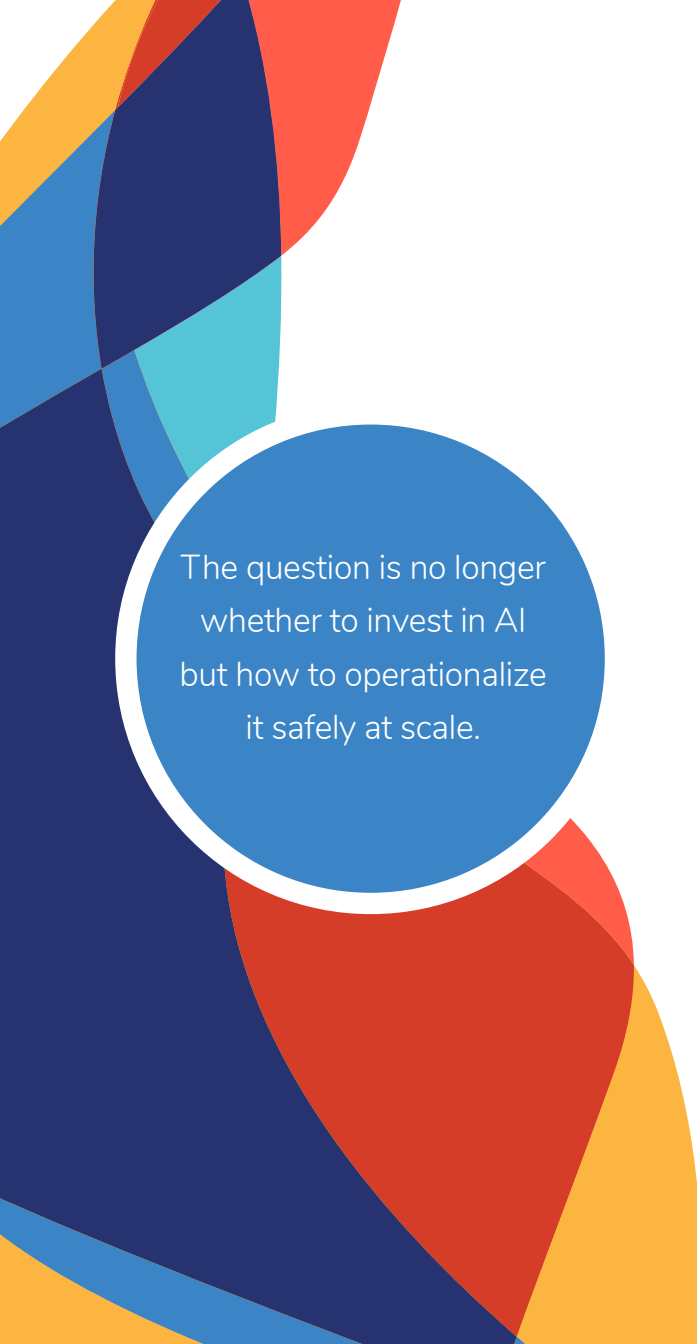
and agents is creating new complexity, prompting leaders to think in terms of a coordinated AI operating fabric rather than isolated tools.

The second wave marks a turning point: enterprises are no longer exploring what AI can do, but building the foundations required to run AI as a dependable, and valuable, part of the business.

Core Priorities for Enterprise AI Maturity Today

As organizations absorb the lessons from their earliest AI initiatives and we move through the second phase, a clearer set of priorities is coming into focus. These priorities represent the foundational work required to move beyond fragmented experimentation and into a stable, scalable phase of AI adoption, especially as [Gartner](#) forecasts the increasing use of agents.





The question is no longer whether to invest in AI but how to operationalize it safely at scale.

- A first priority is consolidating and organizing the information that AI systems rely on. Enterprises have learned that model performance is directly tied to the structure, freshness, and accuracy of their internal knowledge. This is pushing teams to modernize documentation, improve data hygiene, and create governed pathways for AI to access approved content. Without this foundation, even the most capable models struggle to produce consistent, trustworthy results.
- A second priority is establishing common practices for governance. Organizations are standardizing how models and agents are selected, how prompts and workflows are designed, and how AI-driven decisions are monitored and reviewed. Governance is shifting from being perceived as a barrier to being understood as the framework that allows AI to be deployed with confidence across business units.
- Enterprises are also prioritizing the integration of AI into core workflows instead of treating it as a standalone, employee assistant. This means defining where AI should act autonomously, where humans should intervene, and how business rules and policies shape the flow of work. As AI becomes more deeply embedded in operations, this alignment between technology, process, and rules is essential to avoid fragmentation and ensure consistent outcomes.
- Additionally, organizations are investing in the internal capabilities needed to support AI at scale: building new skills, enabling cross-functional collaboration, and fostering a culture that knows how to design, govern, and work alongside AI. These capabilities create the organizational “muscle” required for long-term adoption and prepare the enterprise for the next phase of innovation.


Taken together, these priorities define the present moment for enterprise AI. They form the groundwork for a more orchestrated, enterprise-wide approach in which AI is not a collection of disconnected experiments, but a coordinated, governed layer embedded in the way the organization runs.

Five Predictions for Enterprise AI from 2026 to 2030

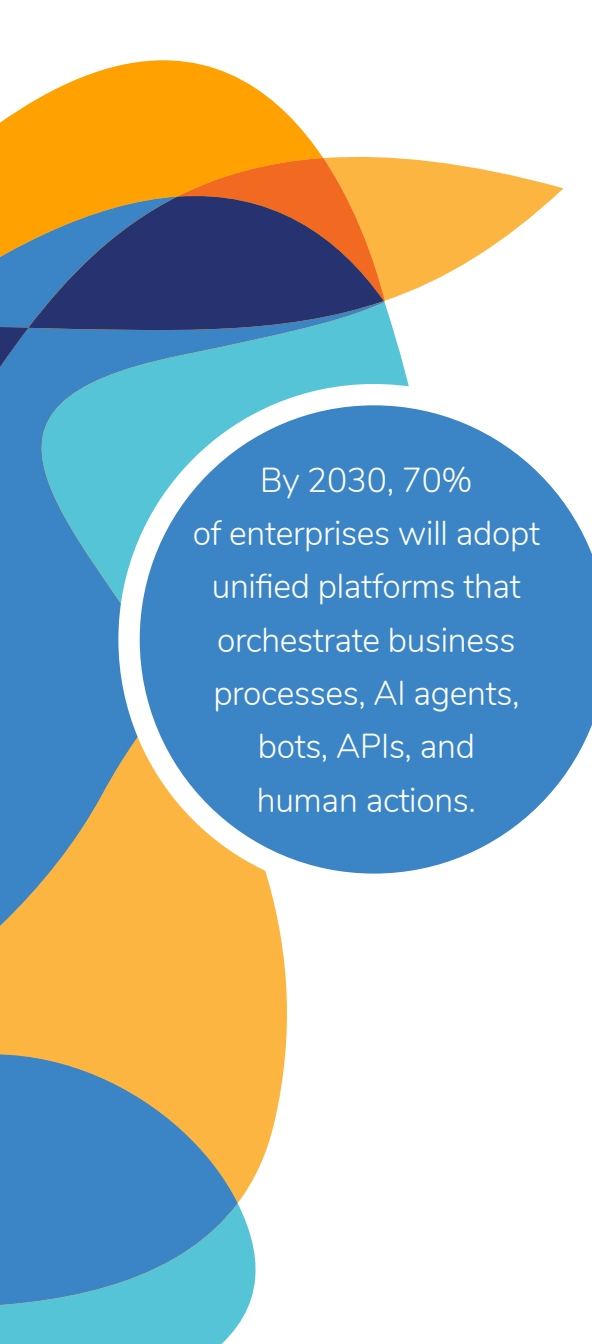
The lessons of the first wave and the priorities emerging today point toward a clear trajectory for enterprise AI. Adoption is no longer driven by curiosity but by the need to modernize systems, elevate decision making, and remove operational friction. As organizations strengthen their foundations, several developments are poised to define the next stage of AI maturity. The following five predictions outline how enterprise AI is expected to evolve over the next five years and what these shifts will mean for technology leaders, operations teams, and the business as a whole.

1 AI ROI will depend more on orchestration than on model selection

Model quality will continue to improve, but gains will be incremental compared to the impact of how models are deployed. The largest performance and ROI gains will come from an organization's ability to orchestrate AI, workflows, rules, and data into a cohesive operating system. Enterprises that invest in a deterministic orchestration



The biggest gains from AI will come from orchestration—not model selection.



By 2030, 70% of enterprises will adopt unified platforms that orchestrate business processes, AI agents, bots, APIs, and human actions.

layer where AI makes decisions, governance turns ideas into validated logic, and execution runs through a rules engine will see faster time-to-value, fewer failures at scale, and a more durable competitive advantage.

2 AI agents will become standard in most enterprises (and welcomed by customers)

Companies will deploy networks of AI agents that support everything from customer service and underwriting to analytics and IT operations. As governance matures, agents will no longer be side experiments but will sit inside orchestrated workflows, with clear boundaries on what they can see, decide, and change. As AI agents become trusted to make decisions and take actions autonomously, they will be embraced by customers to address their needs.

3 BOAT will serve as the central enterprise backbone

BOAT platforms will unify process design, rule management, workflow automation, document handling, and AI orchestration into a single operational layer. Over time, enterprises will expect these capabilities in one backbone platform, using it as the control point for connecting models, agents, people, and systems across the business. And as reported by [Gartner](#), “by 2030, 70% of enterprises will adopt unified platforms that orchestrate business processes, AI agents, bots, APIs and human actions, up from just 5% today.”

4 AI control centers will replace unstructured prompt libraries

Many organizations already struggle with scattered prompts, ad hoc model usage, and inconsistent behaviors. Dedicated control layers will become a standard requirement, managing prompts and policies, monitoring performance, and routing traffic between models. They will orchestrate and govern all aspects of AI use end-to-end.

5 Agentic workflows will replace static flows

Traditional automation relies on deterministic processes, fixed sequences of steps. As agentic capabilities grow, workflows will become more adaptive: AI will propose next actions, surface relevant insights, and dynamically route work to the right person or system. Orchestration and business rules will provide the guardrails, ensuring that these adaptive flows remain compliant, predictable, and aligned with service-level and regulatory requirements.





What This Means for Technology Leaders

For CIOs, CTOs, and heads of operations, these shifts translate into several clear imperatives:

Standardize on an orchestration backbone

Move away from scattered automation and AI tools toward a BOAT platform that unifies workflows, rules, documents, and AI orchestration into a single deterministic control layer.

Treat agents as governed actors, not experiments

Define policies, permissions, and escalation paths for AI agents just as you would for human users and applications and embed them inside orchestrated workflows. Use telemetry and process intelligence to continuously observe how agents behave in production and adjust guardrails accordingly.

Invest in an AI control plane, not just models

Establish centralized capabilities for prompt management, model and agent routing, monitoring, and policy enforcement across all your systems so that every AI interaction is consistent, auditable, and aligned with business objectives.

Make rules and workflows first-class citizens

Use a robust agent rules engine and workflow design to provide deterministic guardrails around non-deterministic models, ensuring repeatable outcomes and clear, granular audit trails in regulated and high-stakes processes.

Build organizational “orchestration literacy”

Develop cross-functional teams that understand processes, data, and AI together, and equip them to design, govern, observe, and continuously improve orchestrated, AI-infused workflows across the enterprise.

Conclusion

The enterprise AI landscape has changed rapidly since the earliest ChatGPT pilots. Organizations that once focused on tools and models are now focusing on governance, orchestration, and long-term operational scale. AI has become a foundational capability that shapes how people work, how decisions are made, and how processes evolve. Future competitive advantages will be determined more by businesses' ability to orchestrate AI, rather than by the underlying AI models they utilize.

This integrated approach gives organizations the ability to automate at scale, modernize key business systems, improve customer experience, and create a foundation for the future of AI-driven innovation.

To learn more about AI orchestration and how it could accelerate your business, **[schedule a time to speak with an AI solution expert.](#)**

