




EXECUTIVE INSIGHTS

An executive's guide to agentic AI

Preparing for the agentic era

Agentic AI is set to be as transformative as the advent of the internet, changing how we organize work, manage operations, and drive value. [Gartner projects](#) that by 2028, one-third of all enterprise software applications will include agentic AI.

Leaders face tremendous opportunities and significant organizational challenges as these AI agents transition from experimental technology to core business infrastructure. This guide provides a concise explanation of what agentic AI is, what makes it unique, the concrete business outcomes it enables, and the practical steps executives can take to prepare their organizations to best take advantage of it.



What is agentic AI?

Agentic AI represents the next evolution in artificial intelligence, moving beyond conversational interfaces to systems that leverage AI to reason, plan, and complete tasks on our behalf.

Instead of the reactive intelligent assistants traditionally associated with generative AI, agentic AI makes use of autonomous agents, which act as proactive virtual collaborators that understand context, iteratively make decisions, and act towards achieving defined goals. They do this by accessing tools, private data stores, and the internet, and by adapting to changing conditions and collaborating with other agents. They remember past experiences, evaluate results, and adjust their approaches to optimize and improve the work they do. They don't just answer questions; they solve problems.

How agentic applications differ from traditional software applications

1

DECOMPOSITION

They're able to take higher-level objectives as inputs and decompose them into plans and any required code necessary to solve the problem.

2

SELF-REFLECTION

Instead of executing pre-defined operations step by step, they're able to dynamically change as they run and iteratively improve.

3

ACTION & TOOL USE

They're able to perform actions in other systems via API calls, access external data sources, and collaborate with other agents that are specialized for particular tasks.

From agents to outcomes

Agentic AI produces tangible business value, significantly enhancing how organizations execute, innovate, and deliver results. Here are three common outcomes agentic AI is enabling today:

Boosting workplace productivity


AI agents automate repetitive tasks and intelligently tackle time-consuming workflows, freeing employees to focus on higher-value activities. Agents can be used to automatically draft documents, autonomously resolve common IT issues like password resets, and speed software development.

To drive a huge boost in developer productivity, [Amazon took an agentic](#)

[approach to legacy application modernization](#).

Using a transformation agent to upgrade production applications from older versions of Java to newer versions, Amazon modernized more than 10,000 applications, saving more than 4,500 years of development time compared to what it would have taken to do manually, realizing \$260 million dollars in annual cost savings from resizing the instances running those Java apps.

Using agentic AI,
Amazon saved more than
4,500 years
of development time
modernizing production
applications.





Rocket Companies enabled faster query resolution, improved personalization, and a better customer experience for those navigating complex home financing processes using an agentic AI-powered support system.

Accelerating business workflows

AI agents streamline critical business processes by reducing task handoff delays and enabling parallel execution. For instance, agents can rapidly perform credit checks, income verification, and document validation simultaneously; or match invoices with purchase orders, validate payments, and quickly resolve discrepancies.

U.S.-based [Rocket Companies developed an AI-powered support system](#) using Amazon

Bedrock Agents, creating an intelligent platform that aggregated 10 petabytes of financial data and provides tailored mortgage recommendations and real-time personalized financial guidance. The results are faster query resolution, improved personalization accuracy, and enhanced customer experience in navigating complex home financing processes.

Speeding innovation and research

Agentic AI can autonomously explore vast datasets, identify novel insights, and generate rapid iterations of ideas. For example, agents can instantly analyze social media, consumer feedback, and market data to identify emerging trends and opportunities; or quickly generate and iterate virtual prototypes based on defined criteria to shrink product development time.

[AstraZeneca, a global biopharmaceutical company, built an agentic solution on AWS](#)


that automates time-consuming manual processes, enabling teams to quickly surface insights across company silos. The system uses autonomous agents that break down complicated tasks into dynamic, multi-step workflows, adapting their approach based on information gathered at each step to deliver comprehensive cross-functional analyses. This agentic solution helps AstraZeneca speed clinical trial decisions, reducing the time it takes to bring new therapies to market while ensuring accuracy, efficiency, and scientific rigor.

AstraZeneca automated time-consuming manual search processes using an agentic solution, accelerating clinical decision making and helping to speed new therapies to market.



Laying the groundwork for agentic AI success


As with any major shift, realizing the full potential of agentic AI will demand deliberate, thoughtful planning. Here are four areas leaders can focus on today to drive agentic success tomorrow.



Double-down on generative AI technology foundations

For many organizations, the technical dimension of agentic AI builds naturally on investments already made for generative AI. Data remains the essential fuel for agentic AI, just as it does for generative AI efforts to date. Enterprises that maintain an AI-ready data infrastructure, complete with unified lakes, vector stores, and

rich metadata, enable the semantic grounding needed to retrieve context accurately and reduce hallucinations. Where a semantic layer has been established, agents can reason across business domains and share insights without repetitive integration work.



Organizations that have operationalized generative AI with production-grade rigor and worked to harness it across the business are likely to realize faster, smoother adoption of agentic workloads.

Organizations that have operationalized generative AI with production-grade rigor and worked to harness it across the business are also likely to realize faster, smoother adoption of agentic workloads. Being able to access the latest models, automatically use the best model for a given task, harness an organization's unique data and apply consistent guardrails, all in a secure cloud environment enabled by tools like Amazon Bedrock, can turn isolated pilots into reusable, governed capabilities that scale quickly across the enterprise while meeting performance, security, and cost objectives.


Collectively, these technical efforts—a modern stack, unified data, semantic context, and operationalized approach—provide a strong foundation on which agentic innovation can scale.

Prepare for the future of human-AI collaboration

One of the most challenging aspects of the agentic era will be managing this change with our people. The rapid rise of generative AI has already sparked trepidation among employees about how their roles will change, and the introduction of autonomous agents will likely continue that trend. Transparent, ongoing communication therefore becomes essential. Leaders who set the context for how agentic AI will be used, which tasks it will solve, and how human expertise will remain central are likely to see faster adoption and higher engagement. Internal forums, open

Q&A sessions, and visible success stories can help build trust with employees and invite them to participate in shaping the new workplace rather than feeling subjected to it.

Once employees understand the “why,” they can more easily embrace the “how.” Frame agents as more than just sophisticated tools. Instead help your teams approach agents as valuable resources that behave like teammates who follow goals, learn from experience, and hand work to one another. Humans retain the irreplaceable



With open communication and targeted upskilling, executives can turn trepidation into enthusiasm and harness agentic AI as a force multiplier of human potential rather than a threat to it.



advantages of moral judgment, empathy, and creative problem-solving, while delegating exhaustive data gathering, statistical pattern recognition, and tireless execution to an agent. New hybrid teams will emerge: a customer-support lead supervising a multi-agent service desk, a business analyst curating insights from research agents, or an operations manager orchestrating supply chain agents that adapt in real time to shifting demand.

This new way of operating demands new skills. Agentic literacy—the ability to supervise, collaborate with, and strategically direct agent teams—will matter as much as traditional digital literacy. Upskilling programs that combine technical know-how with ethical reasoning and change-management skills can help employees move confidently into these expanded positions.


With open communication and targeted upskilling, executives can turn trepidation into enthusiasm and harness agentic AI as a force multiplier of human potential rather than a threat to it.

Embrace flexibility and continuous learning

Agentic AI will reshape more than just workflows—it will challenge the very way companies organize and learn. Where many enterprises still operate as discrete departments, the agentic model resembles an immune system: distributed intelligence that senses issues anywhere and mobilizes resources in real time. For instance, when an agent in finance uncovers a pattern that improves cash-flow forecasting, the insight can be propagated instantly to planning, sales, and supply chain teams—no handoffs required. The value comes from connecting talent, data, and decisions across boundaries rather than optimizing each silo in isolation.

This shift demands that we rethink processes built for linear handoffs. Agents don't wait for a baton; they interpret objectives and orchestrate tasks dynamically as conditions change. Like the adoption of ERP or the move to cloud, leaders will need to reimagine workflows—turning rigid checklists into context-aware playbooks that update continuously. Success metrics, incentive plans, and even org charts may evolve as agent-driven work flows diagonally across functional lines.

The value of agentic AI will come from connecting talent, data and decisions across silos and by embracing discovery and experimentation.





Related to this non-linear approach, is a significant cultural evolution that organizations will need to face: moving from a mindset of perfect execution to one of discovery. Research labs illustrate the power of coupling disciplined methods with openness to unexpected insights, and agent-enabled enterprises need the same balance. Teams should feel free to probe an agent's recommendations, test alternative approaches, and record lessons so both humans and agents improve with every cycle. Rapid feedback loops can allow organizations to learn as much from near misses as from clear wins.

Leaders can set the tone by modeling curiosity: recognizing the value of experiments even when outcomes differ from expectations, rewarding knowledge-sharing and iterative refinements, and spotlighting examples where constructive questioning led to better results. By embracing flexible approaches to agentic work and placing continuous learning above static perfection, organizations can realize the full benefits of agentic AI.


Build a new governance model

Governance and risk management is another area where approaches need to evolve to satisfy a new agentic way of working. Governance in the agentic era looks less like day-to-day supervision and more like a board of directors model. Board members set strategic intent, define success metrics, and specify the decisions that must be escalated—a model suitable for governing agents that enables them to operate independently within defined guardrails. Periodic [“board meetings” for AI](#)—structured reviews that assess outcomes, update policies, and fine-tune objectives—can help keep autonomy aligned with business goals without slipping into micromanagement.

Risk management also needs to evolve. Instead of the factory floor predictability that underpins many traditional controls, agentic AI demands trading floor rules. Just as traders have real-time authority to make decisions within defined parameters, while the firm maintains oversight, agents should have defined risk thresholds they can operate within. Continuous telemetry can help spot drift or correlated behaviors before they compound into systemic exposure, and real-time dashboards, audit trails, and automatic alerting should replace periodic batch reports so that leaders can intervene at machine speed when needed.



Governance in the agentic era looks like a board of directors model: setting strategic intent, defining success metrics, and specifying the decisions that must be escalated.



As agent autonomy increases, so too must ethical oversight. We all expect near-perfect performance from AI—even more than we do from human colleagues—so organizations benefit when decisions are explainable in language that mirrors human reasoning. Accountability never disappears; it redistributes. Clear RACI matrices can spell out who owns data quality, guardrail configuration, and final business approval, ensuring that root-cause analysis is possible when outcomes deviate from an agreed-upon plan.

Privacy boundaries likewise require more than static access controls. Because agents can synthesize context on the fly, they must carry

embedded rules that enforce data minimization, purpose limitation, and jurisdictional constraints in real time. Context-aware guardrails, runtime data masking, and end-to-end traceability help confirm that every decision—no matter how autonomous—remains inside regulatory and corporate policy lines.

Taken together, strategic oversight, real-time risk controls, explicit accountability, and adaptive privacy safeguards form the backbone of responsible agentic AI. By elevating governance from task-level approvals to outcome-level stewardship paired with continuous monitoring, leaders can get the most from agentic AI while still protecting the business.

Start enabling agentic innovation

As leaders, we have a tremendous opportunity to fuel decentralized agentic innovation in our businesses. Consider these 4 steps to successfully introduce agentic AI in your organization:

- **Develop a business-led strategic roadmap for agentic AI implementation.** Start with basic automation and progress toward more autonomous systems with clear governance frameworks and accountability structures.
- **Encourage human-AI collaboration.** Position AI agents as virtual teammates rather than tools, while driving necessary cultural changes and workforce upskilling.
- **Establish security and privacy controls that go beyond traditional methods.** Build in guardrails specifically designed for the dynamic nature of autonomous agents.
- **Finally, decentralize to speed innovation.** Enable AI adoption across departments to drive greater agility and alignment with business needs, only centralizing elements that help contribute to speed such as platforms, tooling, and governance.



The organizations seeing the greatest success today aren't those with the most ambitious plans, but those who have started the learning cycle early, and have gathered real-world feedback to inform their next iteration. Don't try to boil the ocean or wait for all the answers before you begin. Pick a specific business problem that matters and get your teams building. The agentic future is for all of us to create.

Next steps

→ [Hear from your peers](#)

Follow the Executive Insights podcast to see how leaders across industries are approaching the agentic AI future.

→ [Dive deeper on agentic AI](#)

Explore how AWS is helping organizations unlock new value using agentic AI foundations, tools to build AI agents, and agent-powered applications.

