

Unlock the power of agentic AI

An executive guide to next evolution in AI



Table of contents

Introduction

The agentic AI opportunity

3

Chapter 2

Adopt an open platform approach for agentic success

8

Learn more

The future of enterprise automation is autonomous

14

Chapter 1

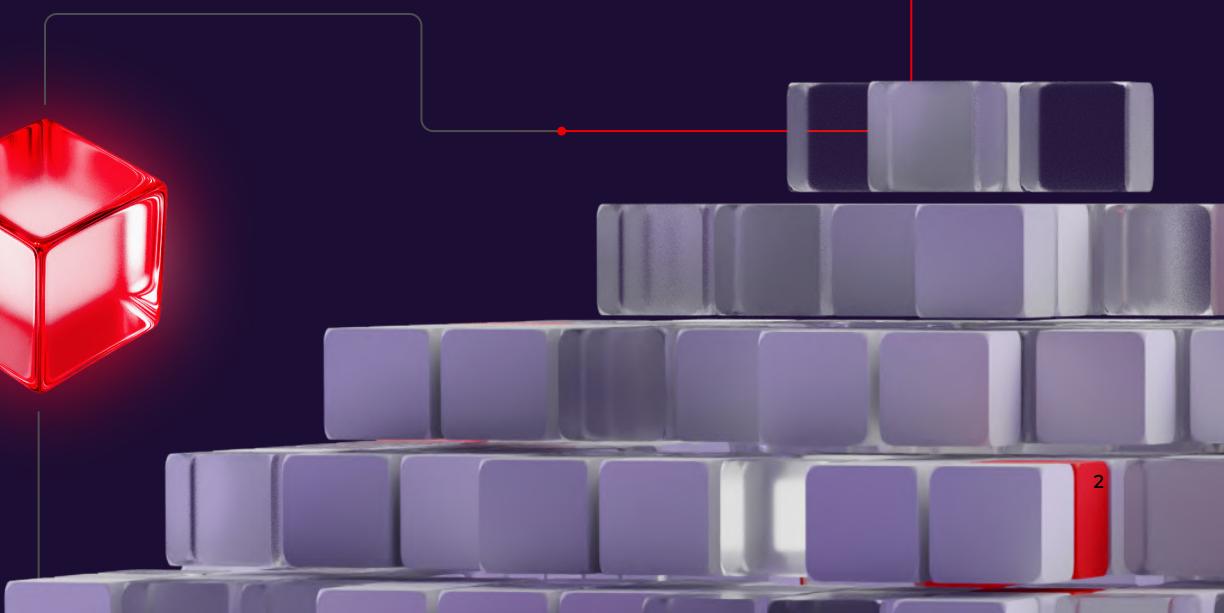
The enterprise's agentic AI challenge

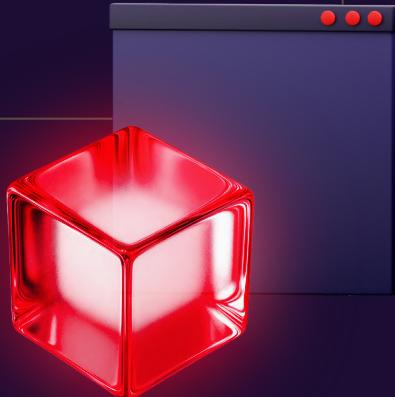
5

Chapter 3

Charting your agentic AI strategy with Red Hat AI

11





The agentic AI opportunity

Today's executives are facing a pivotal moment in the evolution of technology. While AI has already redefined how we interact with data, automate tasks, and serve customers, the next phase in AI is emerging, pushing beyond content generation and predictive analytics.

Agentic AI is an evolution from generative models to autonomous systems with greater adaptability and the capacity for continuous learning. Agentic AI can be a physical structure, a software program, or a combination of both. This allows agents to perceive, decide, and act within predefined parameters across virtually every part of the business. Instead of responding to prompts, agentic AI is capable of initiating multistep tasks, accessing tools and application programming

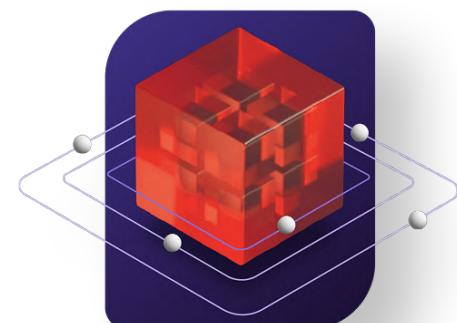
interfaces (APIs), and improving over time. This can look like responding to customer service issues across multiple platforms, automating IT remediation steps, or managing supply chain operations in real time.

For senior leaders, understanding this shift is a strategic imperative for remaining competitive in the future.

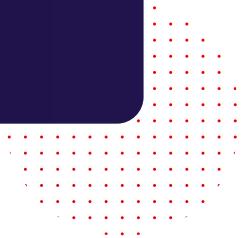
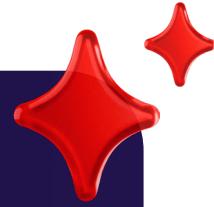
How agentic AI works

Agentic AI is most useful for tasks that require continuous monitoring or rapid decision making. It can be thought of as a way of combining automation with reasoning, decision making, and the creative abilities of a [large language model \(LLM\)](#). To adopt agentic AI into everyday operations, organizations must begin by creating a system to provide an LLM with access to external tools and algorithms that supply instructions for AI agents, which sit on top of other software tools and operate them.

AI agents can then communicate with tools involving orchestration of workflows, depending on the framework being used. This approach allows the LLM to reason and determine the best way to answer a question.



By 2029, agentic AI will autonomously resolve 80% of common customer service issues without human intervention, leading to a 30% reduction in operational costs.¹



The benefits of agentic AI include:



Greater operational efficiency through autonomous workflows that reduce the need for manual intervention.



Reduced costs as a result of minimizing human intervention and improving productivity.



Accelerated decision making and execution across business units, informed by real-time data and contextual insight.



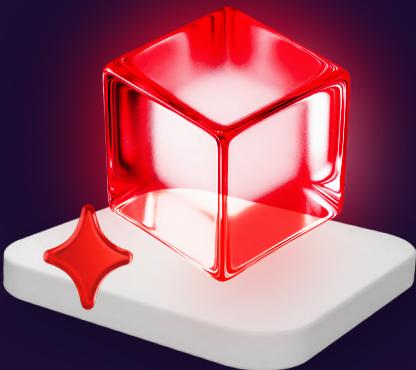
Competitive differentiation, as organizations are able to orchestrate and govern AI agents to increase return on investment (ROI).

In an era defined by constant disruption, agentic AI helps organizations anticipate, adapt, and lead.

The adoption challenge

From orchestrating complex agent workflows, to maintaining trust and reliability, to scaling these innovations efficiently across the enterprise—incorporating agentic AI into your organization will come with challenges. However, with business environments evolving more quickly than ever, adopting agentic AI is no longer optional for organizations seeking to sustain their competitive edge.

This e-book explores these challenges and how to overcome them, the business advantage of agentic AI, and actionable strategies to reshape efficiency and innovation in your organization using agentic AI.



The enterprise's agentic AI challenge

Enterprise knowledge and use of AI is growing quickly, and with that growth has come a level of maturity where more organizations are recognizing its long-term business value.

In fact, AI-enabled workflows—many powered by agentic AI—are poised to expand from 3% in 2024 to 25% by 2026, as investment in AI continues to rise.²

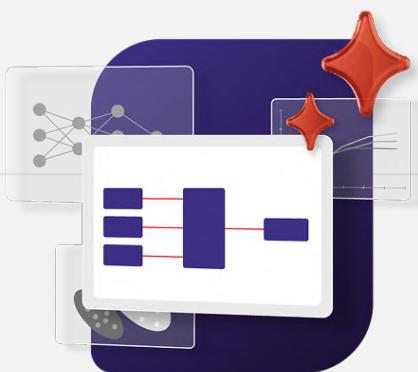
But transformation is never easy. There are real challenges that every organization must navigate before they can deploy intelligent agents at scale and with confidence.

Understanding these challenges early can empower leaders to ask the right questions, design pilot programs effectively, and choose technologies that will set them up for success.

Challenge 1: Orchestrating complex agent workflows

A key differentiator of agentic AI systems, compared to other forms of AI such as gen AI, is that it isn't meant to be a single-point solution.

AI systems can comprehend tools, APIs, communication protocols, data sources, machine learning (ML) models, LLMs, and AI agents. The AI agent is the final decision, but this requires orchestration of the other parts.



² "From AI projects to profits: How agentic AI can sustain financial returns." IBM, 9 June 2025.

This level of orchestration introduces new layers of complexity, including the need to:



Manage across systems

Agentic AI must manage dependencies across systems that can have different protocols, security models, and performance parameters.



Coordinate across agents

Multiagent collaboration must be synchronized carefully to avoid redundant or conflicting actions.



Provide teams the right tools

Developers and IT teams need new tools to design, test, and iterate these workflows quickly so they're not building from scratch every time.

Without a unified framework and a strategy to integrate agentic AI into existing workflows, it risks becoming just another technology in what seems like an ever-expanding digital toolbox.

Challenge 2: Maintaining trust and reliable agent behavior

Before an organization can confidently allow autonomous actions, there must be trust. Unlike other technologies that run with minimal human interaction, such as automation, where outcomes are predictable and predefined, AI agents involve dynamic reasoning and decision making. The core difference is that agents can make decisions autonomously, emulating human behavior, which raises legitimate concerns regarding reliability, transparency, and compliance.

Key questions to consider:



How do you monitor and govern agent behavior in real time?



Can you make sure that agents remain aligned with business policies and regulatory requirements?



What accountability measures are in place to govern and explain an agent's decisions?

Observability, oversight, and proper guardrails are all technical requirements for adoption across stakeholders.

Challenge 3: Scaling agentic AI applications efficiently

Even the most effective pilot can fail at scale. Agentic AI implementations, especially those powered by LLMs, require significant compute resources. Consider the enormous pressure on IT infrastructure running dozens or even hundreds of agents simultaneously. Each agent could be querying models, accessing data, and invoking tools all at the same time.

To scale agentic AI in a sustainable manner, organizations must make sure that:



Their infrastructure has the ability to scale up or down effectively, and that they have methods in place for controlling costs.



They have the flexibility to deploy across hybrid and multicloud environments, closer to data sources or applications.



Governance, security, and performance are all consistent.

Taking a platform approach can address these challenges, allowing organizations to balance innovation with operational best practices.

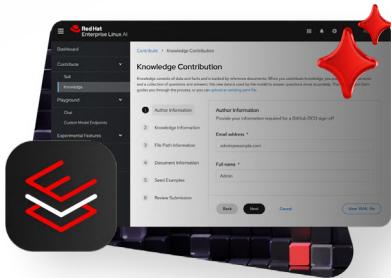


Adopt an open platform approach for **agentic** success

Getting the most from agentic workflows depends on trustworthy infrastructure, operational visibility, security-focused deployment, and flexibility across environments. This is where a platform approach can help.



Built on decades of enterprise open source leadership, [Red Hat® AI](#) is designed to support intelligent agentic workflows, from experimentation to production-scale orchestration. This allows organizations to develop, deploy, and maintain an agentic AI system, including AI agents.



[Red Hat Enterprise Linux® AI](#), for individual Linux server environments.



[Red Hat OpenShift® AI](#), for distributed Kubernetes platform environments.



[Red Hat AI Inference Server](#), for optimized inference of LLMs.

Together, these solutions combine the power of open source technologies with leading open source models, helping organizations accelerate the pace of discovery and democratize access to emerging tools and technologies.

Simplified assembly of agent workflows

Agentic AI success depends on the coordinated operation of planning algorithms, memory stores, tool orchestration, and learning feedback loops. Red Hat reduces the difficulty of building these stacks by offering:



A unified API experience

By standardizing how LLMs, open agent frameworks, and [retrieval-augmented generations \(RAG\)](#) pipelines interact, organizations gain interoperability across diverse model providers and frameworks.



Dedicated AI experiences

Improved user experiences give engineers a single environment to manage and prototype AI assets, agents and applications.



Ecosystem support

Native compatibility with leading LLM frameworks, including LlamaStack, and OpenAI-compatible APIs lets developers choose the right model architecture, including open-weight and fine-tuned options with complex integration overhead.

Adaptable and governed agent deployment

As organizations increase AI agent autonomy, comprehensive governance, observability, and explainability are essential.

Red Hat AI addresses these challenges with features that help maintain responsible deployment:



Context-aware execution

Model context protocol (MCP) allows agents to access semantically relevant operational data and enforce intent-aligned boundaries, improving explainability and control over decision making.



Integrated observability and policy enforcement

Using Red Hat monitoring tools and role-based access controls, teams can enforce behavioral boundaries and trace decisions back to the source models or data.



Interoperable with existing enterprise and compliance policies

Interoperability makes sure that agent activity adheres to organizational security posture across regulated environments.

Red Hat AI empowers enterprises to build autonomous agents that are powerful and predictably aligned to enterprise intent and compliance requirements.



Scalable and cost-efficient AI platform

To scale agentic AI requires more than just big models. Organizations need intelligent infrastructure. This infrastructure must be able to adapt to fluctuating workload needs, optimize cost and performance tradeoffs, and maintain support and security posture.

Red Hat AI provides an approach that offers:



Flexibility across environments

Run agents across the datacenter, cloud environments, clusters, and edge nodes.



Optimized compute orchestration

Support inference cost savings with Red Hat AI Inference Server and reduce total cost of ownership by rebalancing tasks across cloud and on premise resources using Red Hat OpenShift AI to support intelligent autoscaling.



Security-focused scaling

Gain support, security patching, and continuity assurances that are critical for deploying AI in industries such as finance, healthcare, and government.

Red Hat AI is the platform that unifies tooling, governance, and scalability in a single platform. Whether coordinating autonomous agents for IT operations, customer services, or industry-specific tasks, Red Hat meets you where you are on your agentic journey so you can proceed with confidence, not complexity.





Charting your agentic AI strategy with Red Hat AI

With the right foundation, agentic AI can accelerate delivery, reduce costs, and boost ROI. But transformation doesn't happen overnight; it begins with a strategic approach that aligns with business priorities, technical readiness, and governance requirements.

Red Hat AI is built to support organizations at every stage of this journey, from exploration to enterprise-wide deployment. Whether you are just beginning to evaluate agentic workflows or

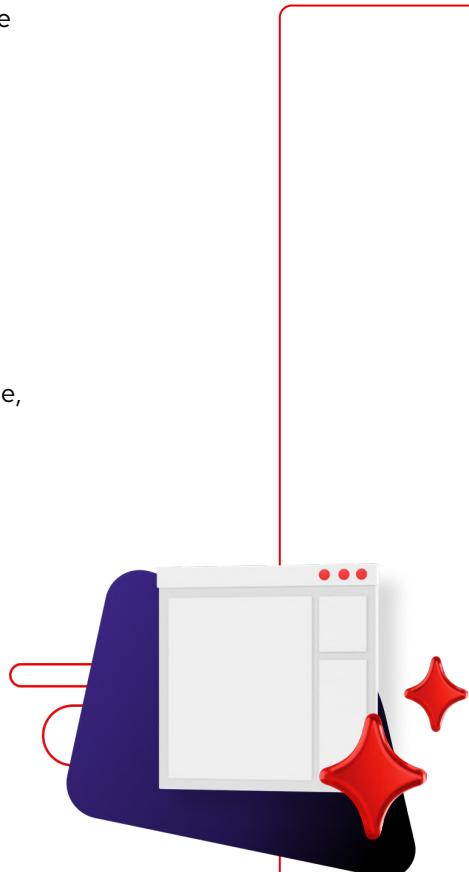
already scaling AI initiatives, a well-planned strategy will make sure that your efforts and investments are focused on generating tangible business value.

Where to start

The path to autonomous operations doesn't need to begin with a full-scale deployment. Rather, a phased adoption that begins by focusing on high-value, well-contained use cases that offer short time to insight is often best.

- Customer support ticket triage and resolution.
- Automated IT remediation or DevOps incident response.
- Compliance report generation across business units.
- Internal knowledge base retrieval agents.

These initial steps provide an ideal testing ground to evaluate agent behavior, establish trust with stakeholders, and validate the performance of your IT infrastructure.



How to assess your organization's readiness

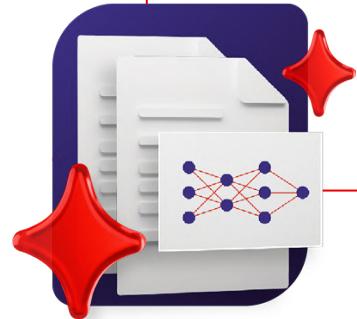
To gain a better understanding of your organization's environment and readiness for agentic adoption, consider the following:

1

Where do repetitive or rules-based decisions slow down my teams?

2

What systems, APIs, or data sources would agents need to access?



3

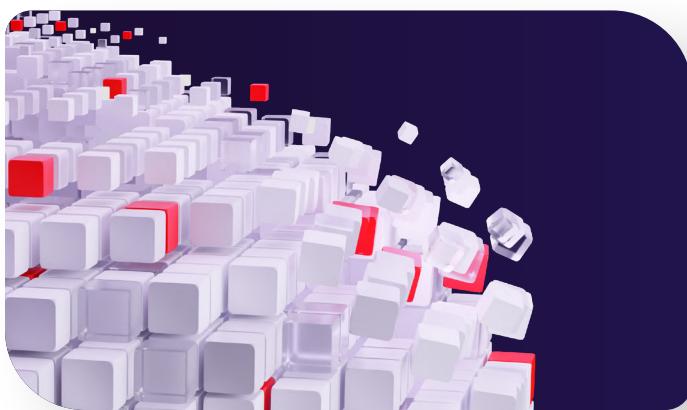
Do I have the IT infrastructure to support low-latency, model-based workflows at scale?

4

What policies and guardrails need to be in place to govern autonomous actions?

5

How will success be measured?
(Consider cost savings, speed, accuracy, and reducing errors in automation.)



While this is not an exhaustive list, these questions are a great starting point for identifying early wins and making sure you have alignment between IT, operations, and business leaders.

Strategic considerations

To set your organization up for success, consider the following elements when you develop your agentic AI strategy:



Interoperability

Choosing an open, modular platform that doesn't lock you into a single model, cloud, or toolset will help you maintain flexibility as technologies evolve.



Governance by design

Establish policies, monitoring, and explainability from the start to maintain compliance and confidence. For example, maintaining compliance with the Health Insurance Portability and Accountability Act (HIPPA).



Embed AI into IT strategy

Treat AI agents as part of your broader digital and automation strategy, not as an isolated experiment.



Skill development

New technology requires new skills. Empower teams across engineering, operations, and business units with training and encouragement to work with intelligent agents.

Agentic AI readiness checklist

With so many moving parts, it can be helpful to develop an internal checklist to assess your organization's position along your agentic AI journey.

- ✓ Have we identified at least 1 business function that would benefit from autonomous agents?
- ✓ Do we understand which systems and APIs an agent will need to interact with?
- ✓ Can our IT infrastructure support model inference workloads with the required governance and security?
- ✓ Do we have buy-in from key stakeholders to pilot agentic workflows?
- ✓ Do we have a plan for measuring success and scaling proven use cases?



[Learn more](#)

The future of enterprise automation is autonomous

Organizations need to see the future of automation not as a means to eliminate manual steps, but as empowering systems to make decisions, adapt to change, and strive for continuous improvement.

Agentic agents can take on complex tasks and act in pursuit of business goals with speed, accuracy, and autonomy.

For executive leaders, this shift signals a new competitive landscape. Those who embrace agentic AI have the potential to unlock new levels of operational agility, innovation, and ROI. And adoption doesn't need to be overwhelming.

Red Hat AI offers an open, hybrid foundation to help organizations succeed. From simplifying agent workflows to governing autonomous behavior and optimizing infrastructure across environments, Red Hat can help forward-thinking leaders operationalize agentic AI in a practical, scalable, and trusted way.

Ready to take the next step in your organization's AI journey?

Learn more about [agentic AI with Red Hat AI](#) or [connect with a Red Hatter](#) to learn more.