orange™ Business

Agentic AI for enterprises

Building autonomy with intent

Editorial

Agentic AI marks a turning point in enterprise technology. It extends beyond output generation to systems that can reason, act, and adapt, introducing new forms of autonomy within business operations. Yet autonomy without intent risks fragmentation and loss of trust. To create value safely, enterprises must design, govern, and adopt these capabilities with purpose.

This white paper brings together four dimensions of that journey. It begins with core concepts, defining what agentic AI is and how to approach autonomy responsibly. It then explores scaling by design, the architectural principles that make agentic systems interoperable, observable, and production-ready.

From there, governance and trust ensure alignment with ethics, compliance, and corporate strategy. Finally, adoption focuses on people: the skills, confidence, and culture that turn technology into transformation.

Together, these perspectives offer a roadmap for organizations moving from experimentation to impact. The goal is simple but demanding: to ensure that as AI systems gain autonomy, enterprises maintain intent, turning intelligence into trusted capability.

In practice, this white paper will help you:

- Decide where agentic AI delivers real value and where human oversight should remain in control
- Build an architecture you can monitor, audit, and scale confidently
- Adopt responsibly, with human-in-the-loop checkpoints and clear, measurable KPIs

Contents

Core concepts: From generation to automation

Enterprises have embraced digital transformation, cloud computing, and AI-driven analytics to stay competitive. Yet much of this focuses on supporting decision-making rather than automation. With the rise of large language models (LLMs) and rapidly evolving AI capabilities, a new frontier has emerged: agentic AI.

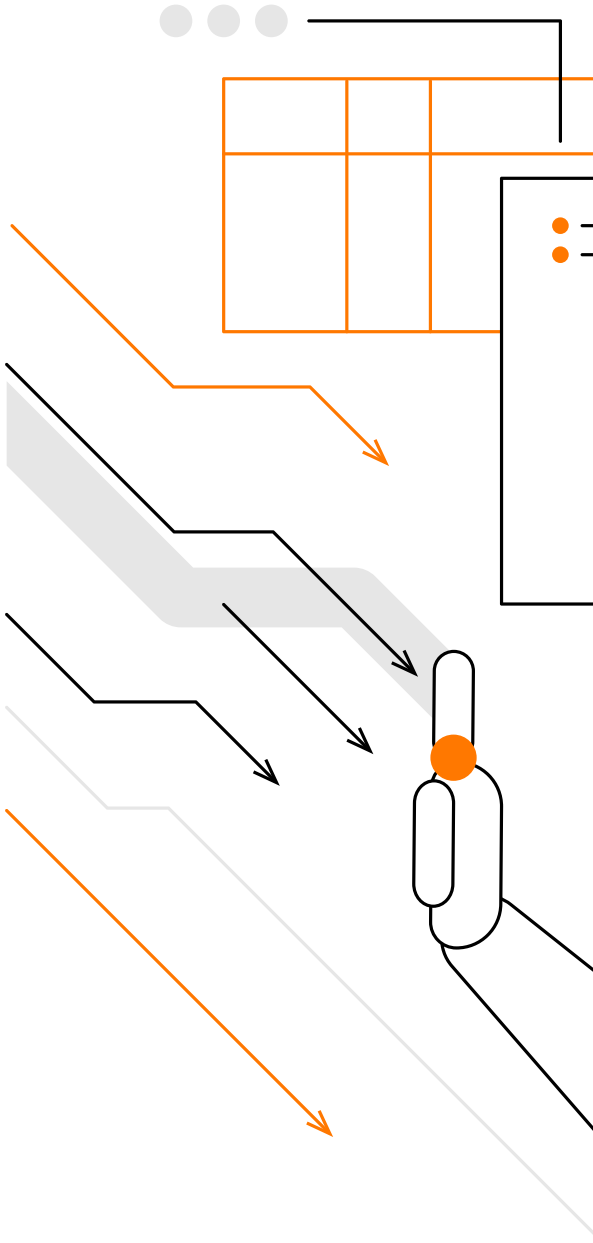

From generating outputs to (semi)autonomous execution of actions.

Agentic AI represents a step-change in enterprise automation, moving from generating outputs (like answering questions) to (semi) autonomous execution of actions to support business processes. This shift offers the possibility of transforming operations, enhancing customer experiences, and empowering employees with intelligent tools.

In this section, we offer a clear view of agentic AI's foundations, defining its core concepts, identifying when it makes sense to apply them, and comparing possible implementation paths. The evolution toward agentic systems opens new opportunities to drive meaningful business impact.

“The evolution toward agentic systems creates new opportunities to drive meaningful business impact.”

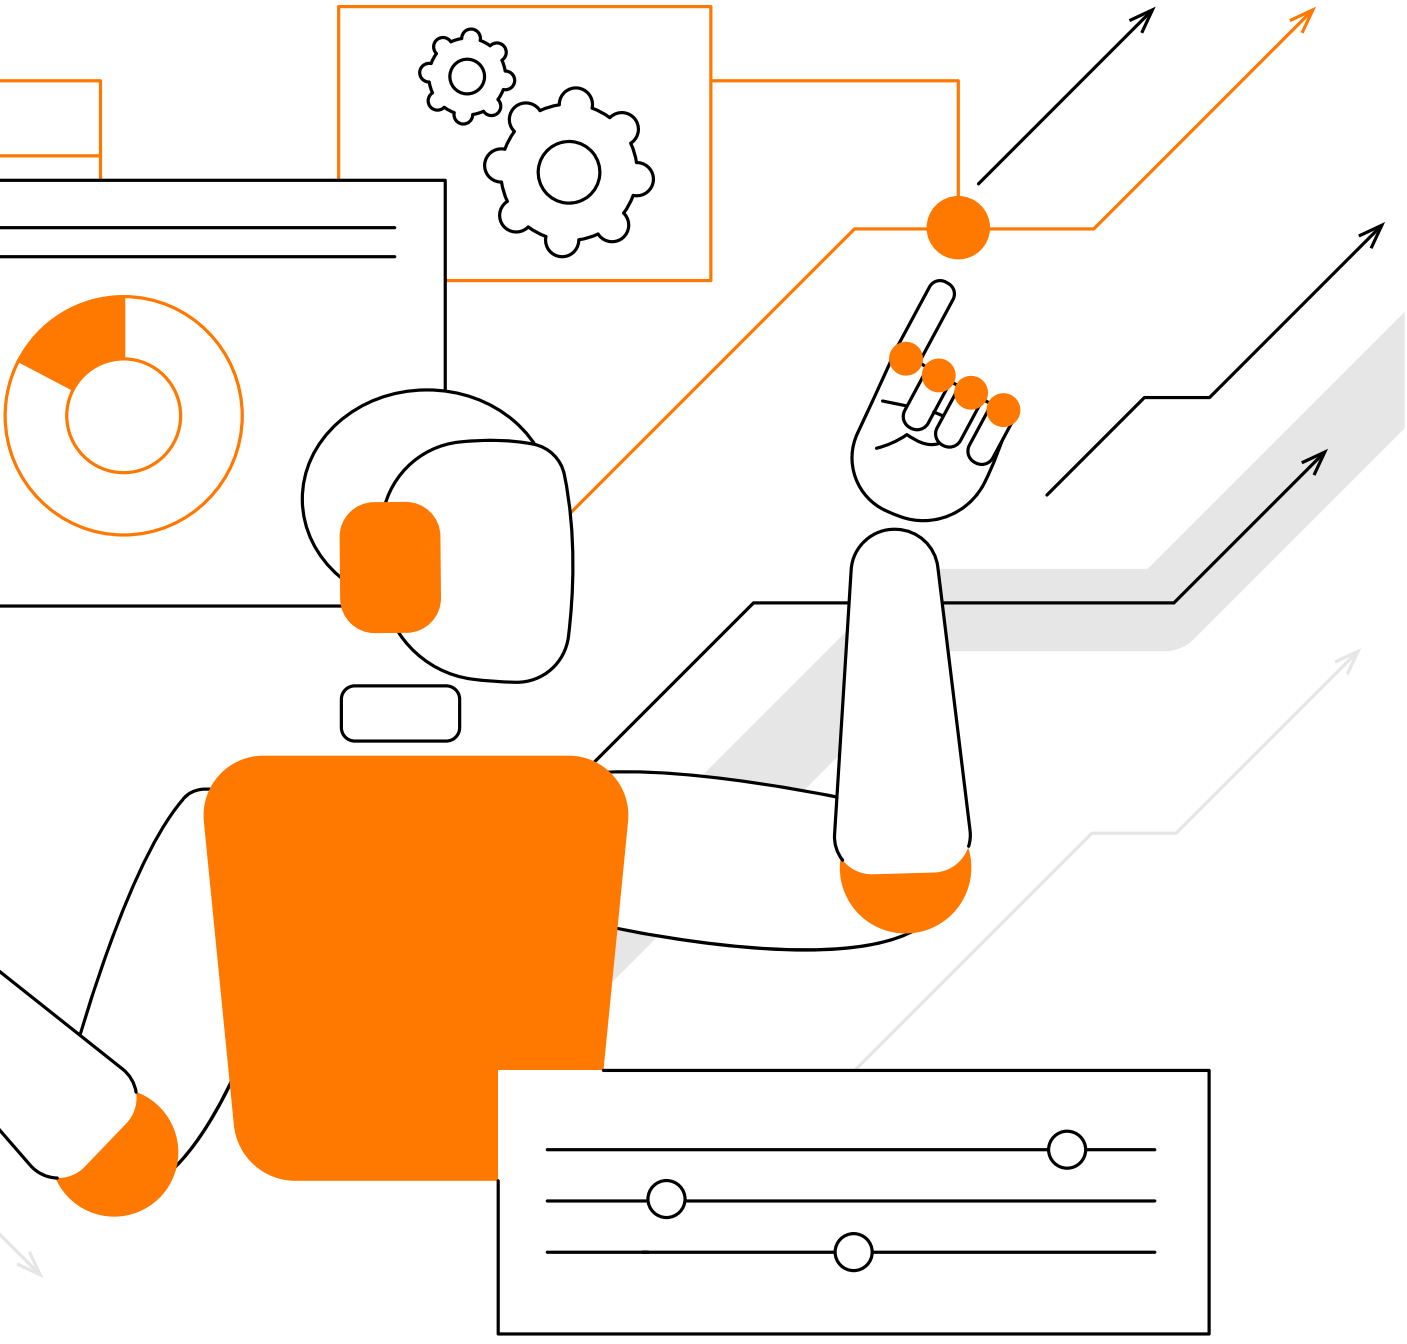

A shift in enterprise AI

Enterprise AI adoption has accelerated in recent years. The turning point came in late 2022, when LLMs such as ChatGPT shifted advanced AI from specialist tools to widely accessible technology, reshaping enterprise workflows and everyday activities.

AI adoption evolution*

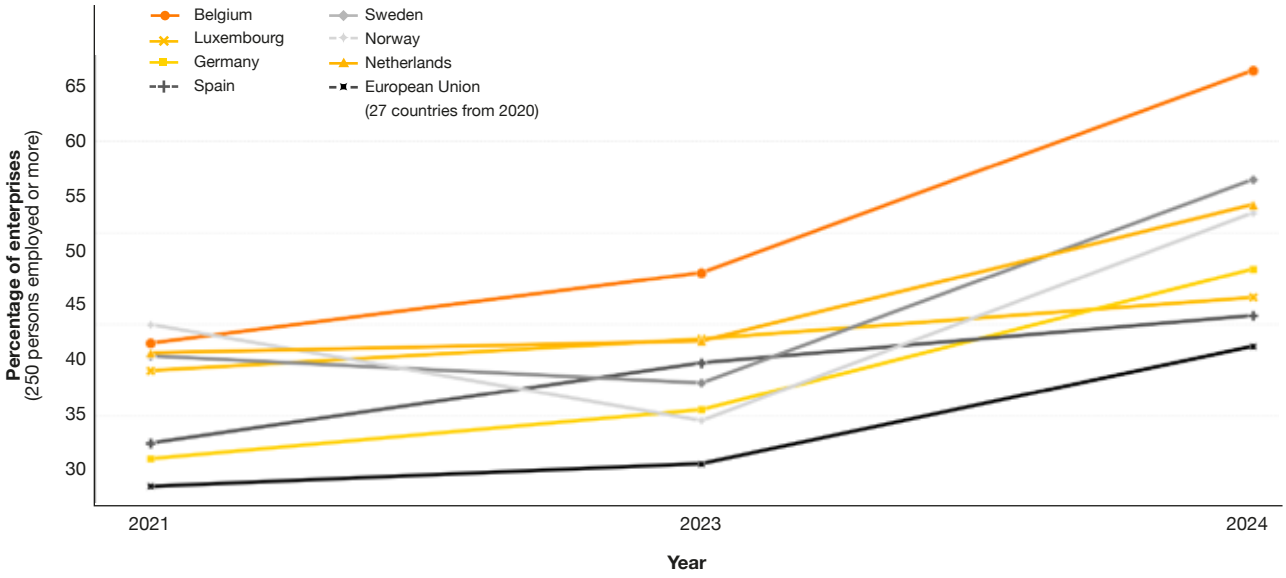

2023 was a year of discovery

Organizations rushed to experiment through Proofs of Concept (PoCs), most centered on tapping proprietary data with Retrieval-Augmented Generation (RAG). These pilots demonstrated potential but also highlighted the limitations of experimentation alone.

By 2024, knowledge assistants powered by LLM-RAG moved into production, supporting core business functions. At the same time, attention turned to agentic AI, as companies recognized that more capable models could automate structured tasks.
Yet rapid evolution has brought both hype and opportunity. New frameworks, methodologies, and vendors are multiplying, leaving many enterprises struggling to find clarity. The challenge is now clear: cut through the noise and build governed, scalable systems that turn agentic AI into real business value.

*Annual AI adoption in enterprises with 250 persons employed or more. Source: Eurostat (isoc_eb_ai).

What is agentic AI?

Agentic AI systems are a class of artificial intelligence designed not only to generate outputs but to act toward defined objectives. They combine reasoning, decision-making, and execution, enabling automation of business processes with or without human intervention.

Depending on autonomy, these systems take different forms:

- **LLM-powered workflows**
  Predefined rules and steps guide the process, for example, a chatbot that retrieves and delivers the right FAQ response.
- **Agents**
  Able to independently reason, decide, and execute multi-step tasks, either as a single agent or as multiple specialized agents working together.

Why it matters: the business case

Agentic AI extends the value of generative AI beyond output creation to process automation. It shifts from answering questions to executing actions that boost productivity and improve experiences.

For example

- Traditional AI might analyze customer feedback and highlight recurring complaints.
- Agentic AI goes further, automatically alerting the product team, drafting next steps, and even executing fixes or campaigns.

Key benefits

Efficiency

Automates routine tasks, freeing up time for higher-value work.

Faster Resolution

Enables real-time monitoring, triage, and escalation.

Cost Savings

Reduces manual interventions and operational overhead.

Personalization

Powers proactive, tailored recommendations at scale.

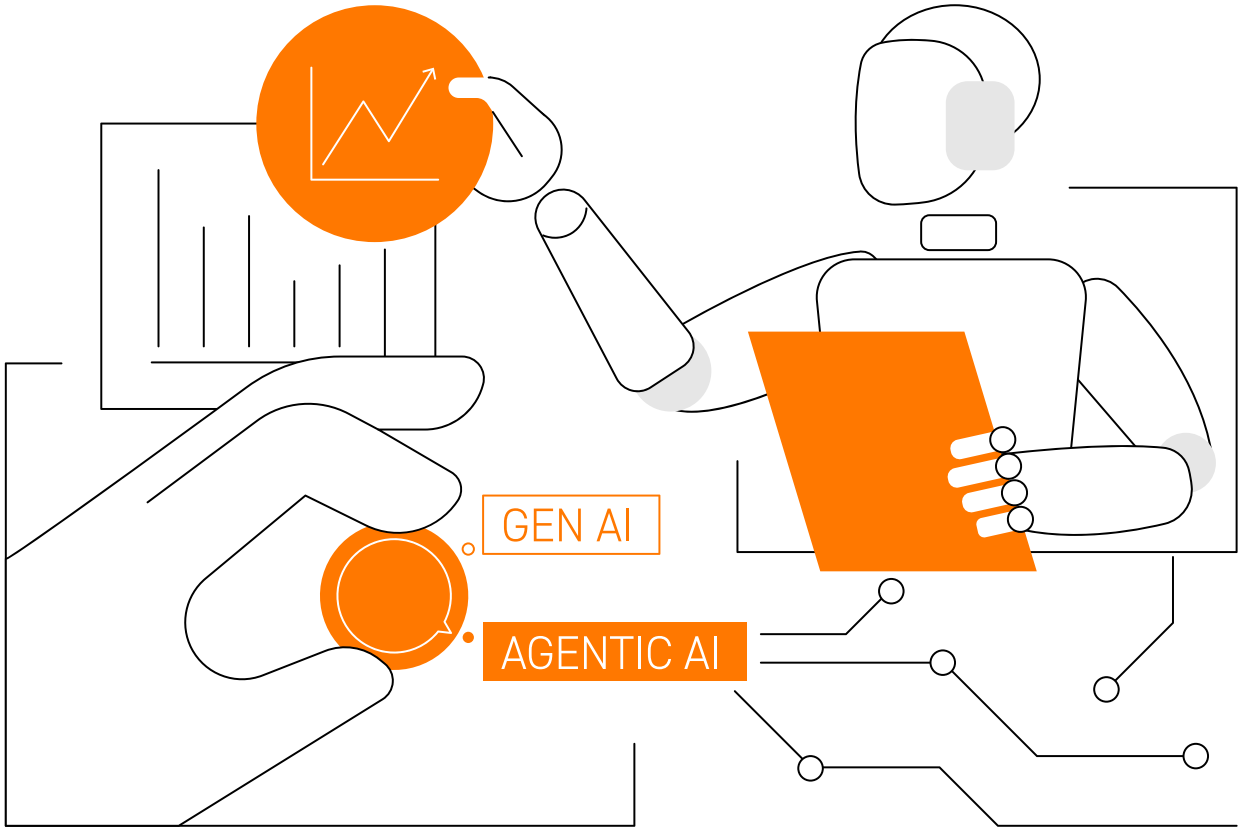

Core building blocks

Agentic AI systems are built from interoperable components:

Large Language Models (LLMs)

Provide reasoning and natural language understanding. Examples: Mistral, Gemini, GPT.

Tools

Extend the system's reach into applications and external data sources, from knowledge retrieval to task execution. Tools may themselves be agents.

Prompts

Define objectives, guardrails, and behaviors through structured instructions and context.

Choosing the right system

Agentic AI is not about deploying agents everywhere. The appropriate level of automation depends on the complexity of the process.

Workflows

Are best for rule-based tasks, where objectives can be expressed in simple, predefined steps. They rely on clear inputs and outputs, making them easier to govern, audit, and scale.

Single agents

Fit dynamic processes that require contextual reasoning and adaptive decision-making. Useful when the complexity of a process cannot be captured in fixed rules. The AI model drives the decision process, determining each next step, iterating through multiple actions, and working toward the outcome.

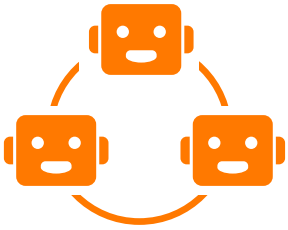

Multi-agent systems

Suit highly complex scenarios where specialized agents must collaborate. These agents handle part of the process, passing inputs and outputs between them until the objective is achieved. This design is powerful for orchestrating interdependent tasks across domains, but also harder to govern and trace.

Feature	Workflow	Single Agent	Multi-Agent
Autonomy	Low	Medium to High	High
Latency	Low	Medium to High	High
Flexibility	Low	High	High
Governance Complexity	Low	Medium	High
Cost / Resource Usage	Low	Medium to High	High
Traceability	High	Medium	Low

The trade-offs grow with autonomy. Agents generally introduce higher latency, increased cost, and greater governance complexity due to their iterative and non-deterministic nature. With multi-agent architectures, the challenges intensify: debugging, monitoring, governing, and scaling become significantly more difficult.

Looking ahead, these principles provide the foundation for building scalable agentic AI systems. The next section explores how to translate intent into design, optimizing key components such as LLMs, tools, and prompts to create modular, production-ready architectures that drive scalability, reliability, and impact.

Guiding Principles

Agentic AI is moving rapidly from proof of concept to production. Its growing capabilities open new opportunities to enhance employee productivity and automate business processes. However, success comes not from adopting the most advanced architecture, but from matching the right level of autonomy to the problem.

- **Start simple with workflows for structured tasks.** Many objectives can be met with well-structured workflows, where tasks are predefined and executed as a clear sequence of steps.
- **Adopt single agents** where reasoning and adaptability are needed, such as tasks that are less predictable and require contextual reasoning or adaptive decision-making.
- **Move to multi-agent orchestration** only when specialization and collaboration become critical, for example, when single agents begin to mismanage tasks or select inappropriate tools.

Modularity is key as sophistication should grow gradually, grounded in proven business value. By aligning architecture choices to process complexity and end-user requirements, enterprises can unlock the benefits of Agentic AI while keeping systems efficient, governed, and future ready.

Your next steps — starter plan

1 Start with clarity

Define the business goals, use cases, and context where agentic AI can deliver measurable value, before building anything.

2 Choose autonomy with intent

- Use workflows for predictable, deterministic tasks.
- Use agents only for reasoning-heavy or adaptive steps.

Always aim for the smallest sufficient autonomy that achieves the business outcome.

3 Ensure human oversight

Keep humans in the loop to guide, supervise, and adjust the system as it operates.

Part 02

Scaling by design

Building production-ready agentic systems

Agentic AI represents a significant evolution of AI, enabling systems not only to generate content, but to act, “reason”, and adapt based on context. This shift allows AI to move beyond passive output generation toward active decision-making and participation in business processes.

Enterprises are increasingly exploring the potential of AI to enhance productivity. Yet, its broader economic impact remains uncertain, particularly when integrating these systems into existing operations. Many organizations find themselves stuck in pilot mode, struggling to operationalize proofs of concept and embed them within their core business functions.

The difficulty lies less in the technology itself and more in the robustness of how the systems are built, governed, and aligned. Transitioning to production demands attention to several key considerations:

- **Strategic alignment:** ensuring the system directly supports measurable business objectives.
- **Governance and risk management:** defining oversight, compliance, and accountability mechanisms.
- **Data readiness and integration:** preparing, connecting, and securing data that drives agentic behavior.
- **Technical design and scalability:** building modular, resilient architectures that can evolve with enterprise needs, and the evolution of AI capabilities.

In this section, we delve into the technical design dimension, exploring how to architect scalable, production-grade agentic systems that can be embedded into core business functions and deliver measurable value.

Building scalable agentic systems

Agentic systems are modular by design, composed of interoperable components that meet specific needs. Building for production requires deliberate choices about these components and how they interact.

The key, though, is finding the sweet spot between:

- Do not overengineer. Excessive complexity slows progress and limits adaptability.
- But avoid “quick-and-dirty” solutions. Designs that bypass production-grade requirements create fragility and rework later.

A robust design in practice means:

- Selecting appropriate models aligned with the problem domain and operational constraints.
- Defining data pipelines and integrations that ensure reliable access to high-quality, contextual data.
- Providing models with the right context and clear, well-structured instructions.
- Establishing continuous monitoring and feedback loops early in the prototyping phase, allowing teams to incrementally add/update features or components, strengthen security, and improve resilience as the system and AI evolve.

How to choose the right model

There is no one-size-fits-all model. Selecting the right model depends on a clear understanding of the business problem, data landscape, desired outcomes, and model limitations. Models vary significantly in their capabilities, latency, and cost-efficiency.

A pragmatic approach involves an iterative and evidence-based evaluation process:

- Prototype with the most capable model to establish a performance and quality baseline.
- Experiment with smaller or specialized models to determine whether comparable results can be achieved at lower cost, latency, and carbon footprint.
- Adopt a hybrid model strategy where tasks of varying complexity are distributed intelligently (Highly capable models -> reasoning-intensive tasks, Lightweight models -> routine or narrow-scope functions).

“The key is to design for seamless model substitution and orchestration while addressing data sensitivity and compliance requirements.”

The key is to design for seamless model substitution and orchestration while addressing data sensitivity and compliance requirements. In parallel, establish a consistent evaluation framework to benchmark model performance. Finally, ensure operational resilience for model decommissioning and version transitions. Together, these practices enable teams to evolve systems as model capabilities, pricing, and data availability change, without major reengineering.

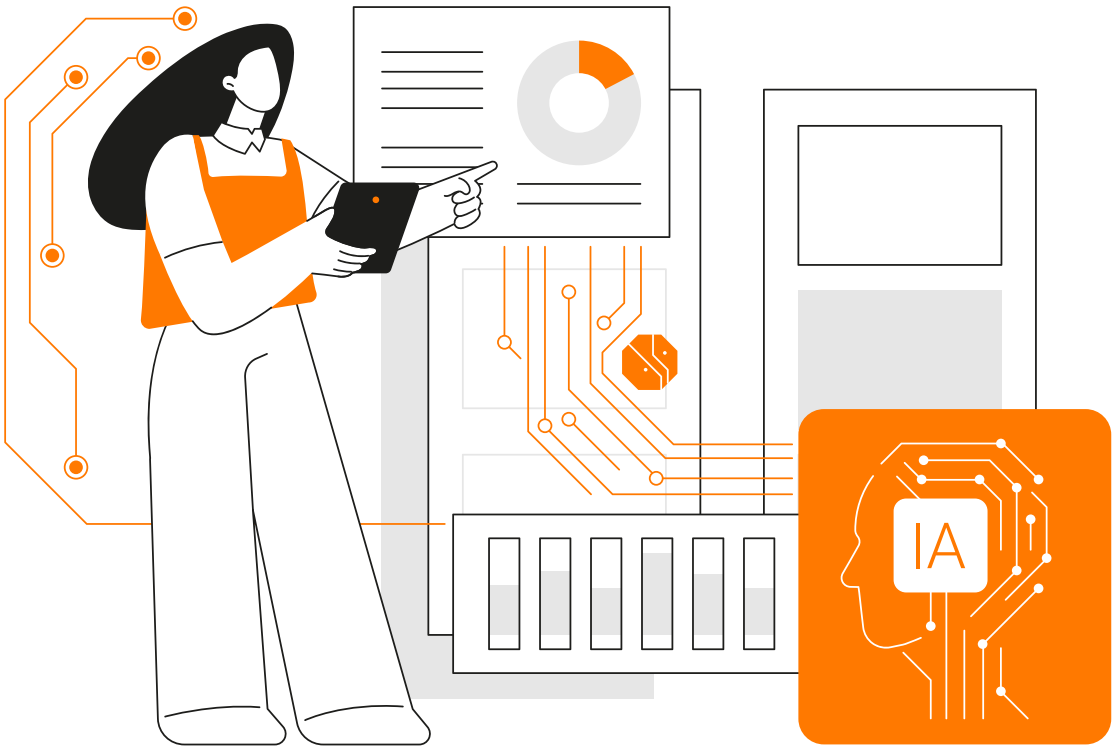

State of AI

- Inference cost keeps declining.
- Small models are getting more capable, making inference faster and cheaper (e.g., Microsoft Phi family).
- The gap is narrowing between open (e.g., DeepSeek) and closed models (e.g., OpenAI GPT-4o)—and even between frontier models themselves.
- Task performance is improving (text, images, video, code). Complex reasoning and planning remain challenging. Reasoning-first models (e.g., OpenAI o1, o3) use chain-of-thought–style inference, hence better reasoning, but at the trade-off of higher latency and cost.

LLM inference prices have fallen 9x to 900x/year, depending on the task

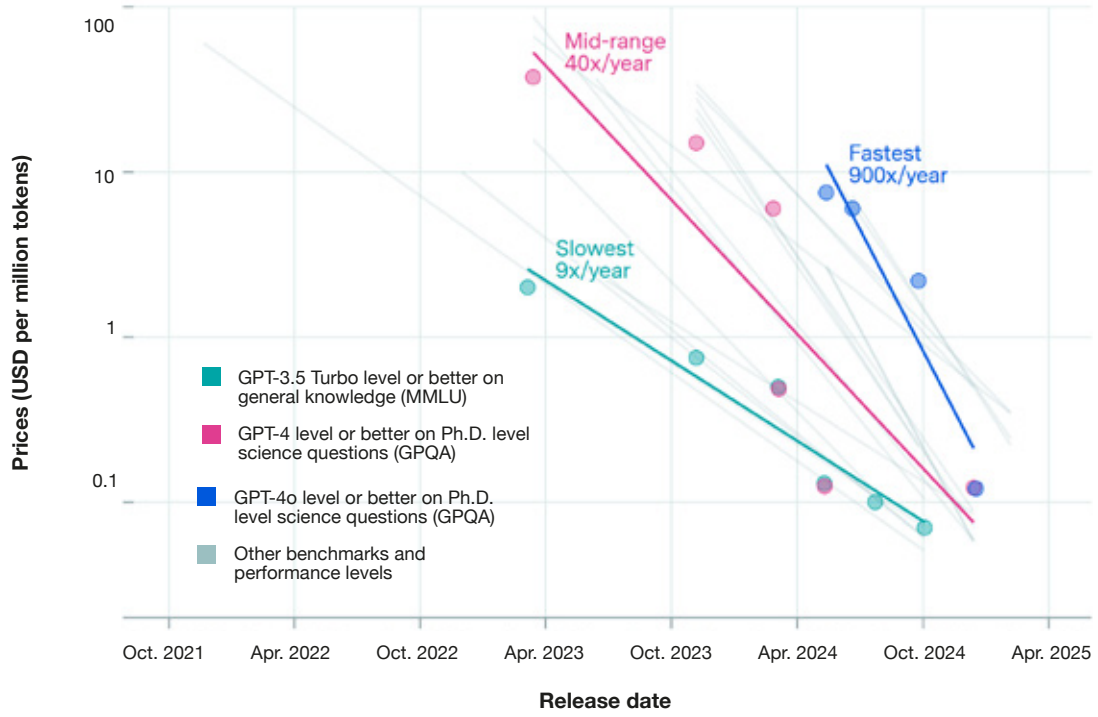

Source: Ben Cottier et al. (2025), "LLM inference prices have fallen rapidly but unequally across tasks". Published online at epoch.ai. Retrieved from: 'https://epoch.ai/data-insights/llm-inference-price"

Open vs Closed

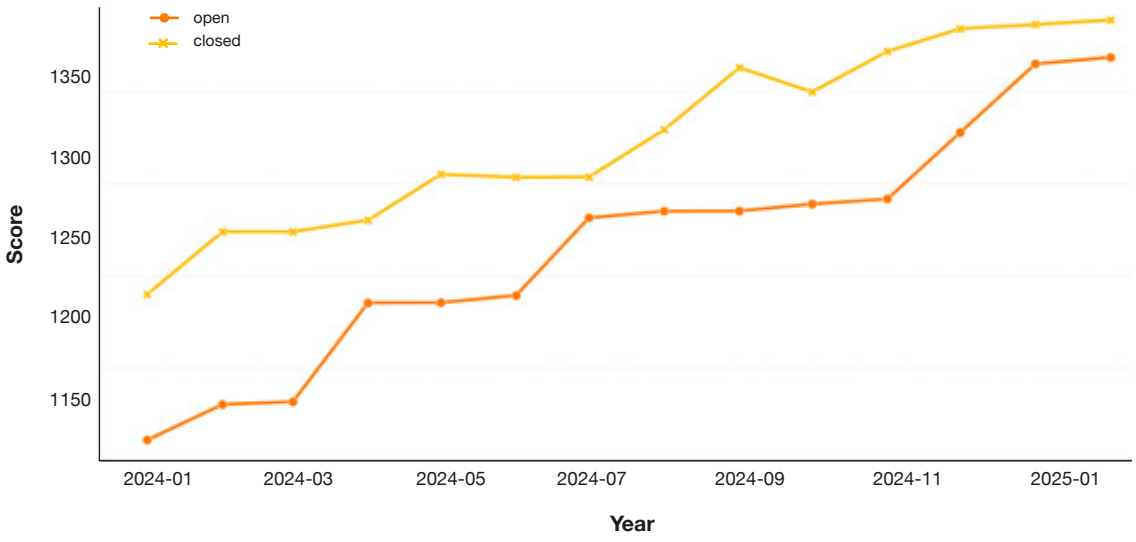

Performance of top closed vs. open models on LMSYS Chatbot Arena. Source: Maslej, N., et al. (2025). The AI Index 2025 Annual Report. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.

Beyond LLMs

LLMs alone barely deliver business value. The true benefits emerge when they are connected to external data sources, systems, and processes.

Tools, actions and protocols

Achieving value requires extending LLMs with the ability to access context, take actions, and collaborate. From the field, some lessons stand out:

1 AI-ready data is a strategic differentiator

It is about data quality together with governance, accessibility, and multi-modal coverage across text, image, and audio formats. Models become valuable when they can consume and act on this data through well-defined and secured integrations.

2 Security and governance ensure trust

As LLMs gain access to tools and data, enforcing authentication, authorization, Human-in-the-loop (HITL), fallback mechanisms, and auditability becomes essential. Strong governance and observability safeguard compliance, prevent misuse, and maintain accountability.

3 Integration unlocks capability

The strength of an agentic system depends on how seamlessly it can access external context and tools. Robust integration mechanisms enable systems to reason, act, and adapt.

4 Interoperability is accelerating innovation

As the ecosystem matures, emerging standards for tool and agent coordination are reducing the need for custom extensions.

Building beyond LLMs is not about replacing models with agents; it is about connecting and orchestrating intelligence. Success depends on how well systems integrate data, orchestrate capabilities, and adopt emerging standards that make agentic collaboration scalable, secure, and sustainable.

“Achieving value requires extending LLMs with the ability to access context, take actions, and collaborate.”

Key developments

Anthropic’s Model Context Protocol (MCP):
Defines how agents connect to external systems — both for context and capability. It provides standardized access to tools, resources, and data sources, effectively extending what a model can “see” and “do.”

Google’s Agent-to-Agent (A2A)
and IBM’s Agent Communication Protocol (ACP): Structure agent-to-agent communication. These standards enable agents to discover, exchange, and coordinate dynamically, allowing them to collaborate, share data, and distribute work efficiently.

Context engineering

Context is the operational fuel of LLMs. It defines how models perceive their environment, the information they receive from systems, users, and external sources.

- Prompts serve as the control interface, defining objectives and guiding behavior.
- Context engineering ensures the right information flows at each turn, complementing prompt design by keeping inputs relevant, current, and scoped appropriately.

Poorly designed prompts or unmanaged context lead models to overgeneralize, misinterpret ambiguous input, or miss critical signals, resulting in poor results or unsafe behavior. Sustained reliability in production comes from treating prompt and context design as iterative engineering practices, continuously tested, versioned, and refined as systems evolve.

“Context engineering ensures the right information flows at each turn”

Some strategies

- Use simple, clear instructions, precise enough to guide behavior, but not overly restrictive.
- Organize prompts in logical sections: background, instructions, tool usage, and output formatting.
- Iteratively refine and version prompts. Continuously improve prompts by adding targeted examples and instructions based on observed failures. Treat them as evolving artifacts, benchmark performance, monitor drift, and maintain version control to ensure long-term reliability.
- Ensure tool definitions are unambiguous and self-contained, with clear purpose and relevant outputs.
- Apply few-shot prompting to provide sufficient examples without overloading with edge cases.
- Leverage metadata (e.g., file size, naming conventions, timestamps) as signals for relevance and priority.
- Use standardized templates to ensure consistency across agents and use cases.
- Include fallback guidance to handle ambiguous or out-of-scope queries gracefully.
- Adopt long-context strategies:
  - **Compaction:** periodically summarize and restart the context with key information.
  - **Agentic memory:** enable agents to record and retrieve essential notes.
  - **Sub-agent architectures:** delegate specialized tasks to focused agents.

Orchestration

Orchestration coordinates and manages how system components work together toward a defined goal. The required level of orchestration depends on system complexity, shaped by the number of tools, agents, and decision points involved.

Effective orchestration is not about building the most complex network of agents. Exhaust the capabilities of simpler systems before introducing additional layers of orchestration. Each new agent or tool adds flexibility but also overhead in governance, monitoring, and troubleshooting. The most successful systems evolve from clear, well-instrumented workflows into modular, multi-agent architectures guided by operational data and business value.

Key considerations

✓ Start with simple workflows

For low-complexity use cases, begin with single-agent systems or orchestrated workflows. They provide faster deployment, easier maintenance, and lower operational risk.

✓ Scale to multi-agent architectures only when necessary

Transition to multi-agent systems when simpler designs can no longer handle tool selection, task decomposition, or coordination effectively.

✓ Design for observability, security, and governance

Ensure end-to-end visibility into task execution, dependencies, outcomes, agent identities, and cost.

✓ Define routing and planning policy

Start with deterministic rules; introduce learned routers/planners as complexity increases.

✓ Choose a coordination pattern

Be explicit about whether you use a central orchestrator or a hierarchical planner/executor/verifier, and document escalation paths and HITL checkpoints.

Some multi-agent coordination patterns

Pattern	Flow	When	Example
Sequential orchestration		Tasks must be completed in a fixed, linear order. Each agent passes its result to the next.	**Document processing** Agent A extracts text, Agent B analyzes, Agent C generates a summary.
Concurrent orchestration	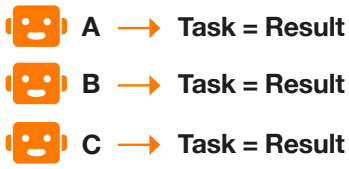	Multiple agents work in parallel on different tasks or the same task, providing independent results.	**Parallel search** Agents A, B, and C simultaneously search different databases for relevant articles.
Group chat orchestration	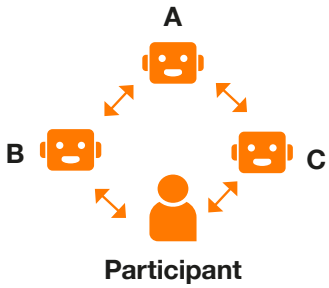	Agents and humans collaborate in real-time, discussing and refining ideas to reach a solution.	**Creative brainstorming** Agents and humans discuss ideas to generate innovative solutions for a marketing campaign.
Handoff orchestration	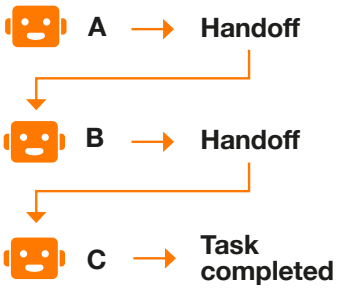	Tasks are passed between agents based on expertise or the agent's ability to continue.	**Customer support** Agent A starts a customer query, hands it off to Agent B for technical details, and then Agent C resolves the issue.
Magentic orchestration	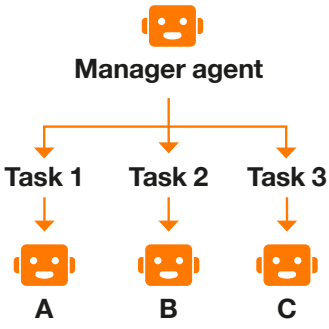	The manager dynamically assigns tasks based on changing needs and new information.	**Project management** A manager assigns specific tasks to agents as the project evolves, adjusting based on progress and new requirements.

Observability and evaluation

Evaluating GenAI differs from Machine Learning (ML): outputs can vary between runs and depend on the context provided. The aim is to determine how well the system meets its intended objectives across the lifecycle.

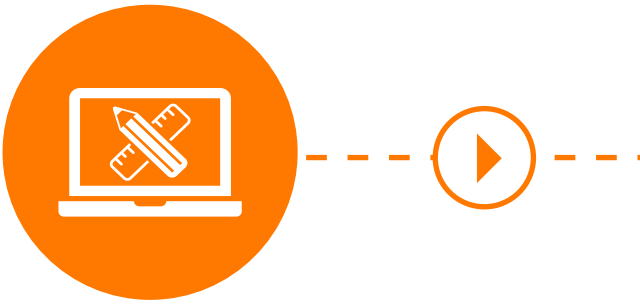

Design

- Define success in terms of business goals, e.g., reduced response time, fewer compliance violations, faster end-to-end processes, or higher user satisfaction.
- Use test sets (golden/edge/stress).
- Plan what to track: steps completed, response time, costs, quality signals.
- Decide who review what and how feedback is captured (e.g., thumbs up/down, issue list).
- Set guardrails: privacy, brand/safety policies, and data handling rules.

Key-metrics:

- RAG (retrieval, groundedness, response completeness)
- Agent-specific (intent resolution, tool call accuracy, task adherence)
- Risk & Safety (Code vulnerability, violence, self-harm)

Considerations

In practice, three considerations keep efforts aligned and measurable:

1. **Evaluation framework**
   Define how you compare design choices and map technical metrics to business KPIs so there is a direct line to strategic goals.

2. **Continuous monitoring**
   Track performance and compliance in production, watching for drift, latency/cost spikes, and reliability issues.

3. **Human-in-the-loop (HITL)**
   Use structured reviews for approvals and to capture feedback that feeds learning loops and ongoing improvements.

Treat observability and evaluation as a built-in capability. Systems stay reliable when metrics tie to business goals, traces make behavior explainable, and HITL feedback closes the loop.

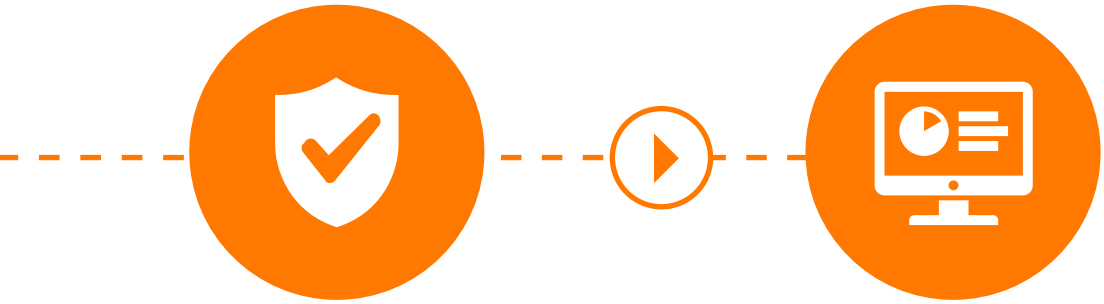

Pre-production

- Trial on realistic cases, including edge and stress scenarios.
- Run safety checks: privacy and policy adherence.
- Shadow test on real traffic.
- Dry-run dashboards and alerts.

In-production

- Live dashboards & alerts on quality, safety, reliability, speed, and cost.
- Explainability: being able to explain “what happened and why” for any case.
- Watch for drift and trigger re-checks when results or inputs change.
- Runtime safeguards: automatic fallbacks, staged rollouts, feature flags.
- Regular audits and reviews, manage capacity and keep costs within budget.

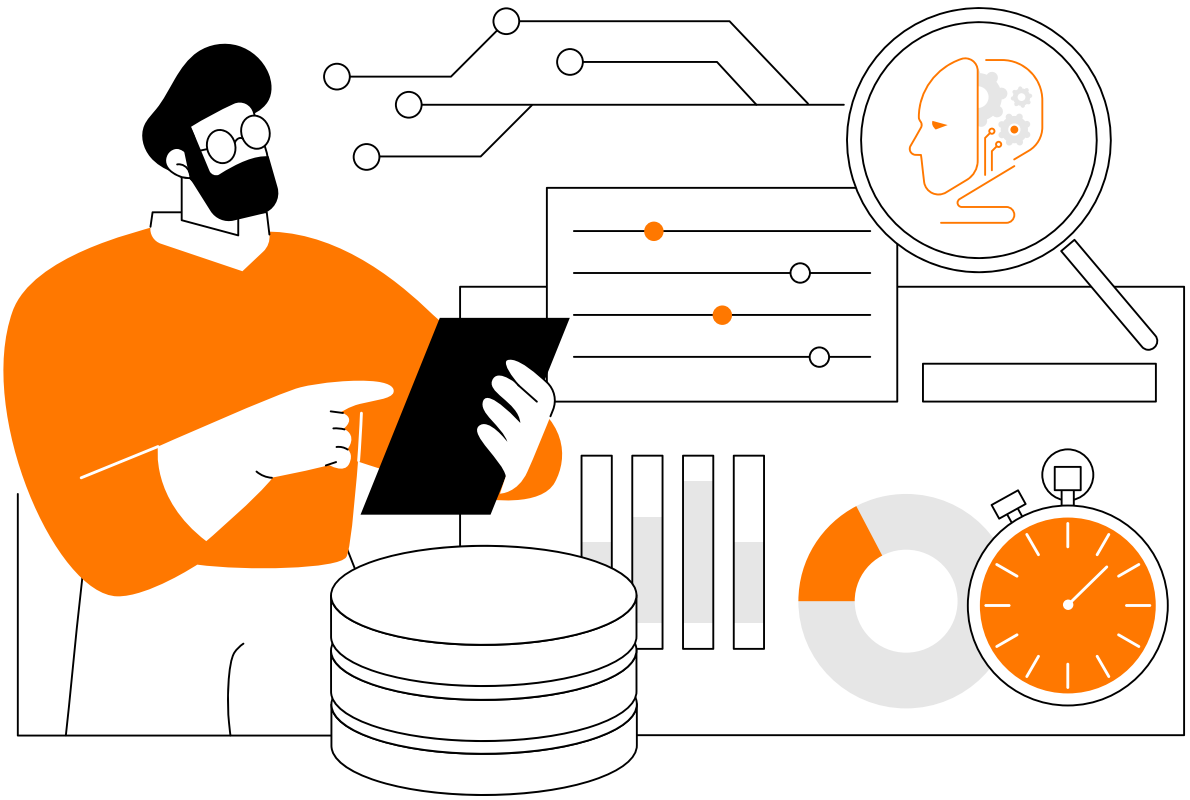

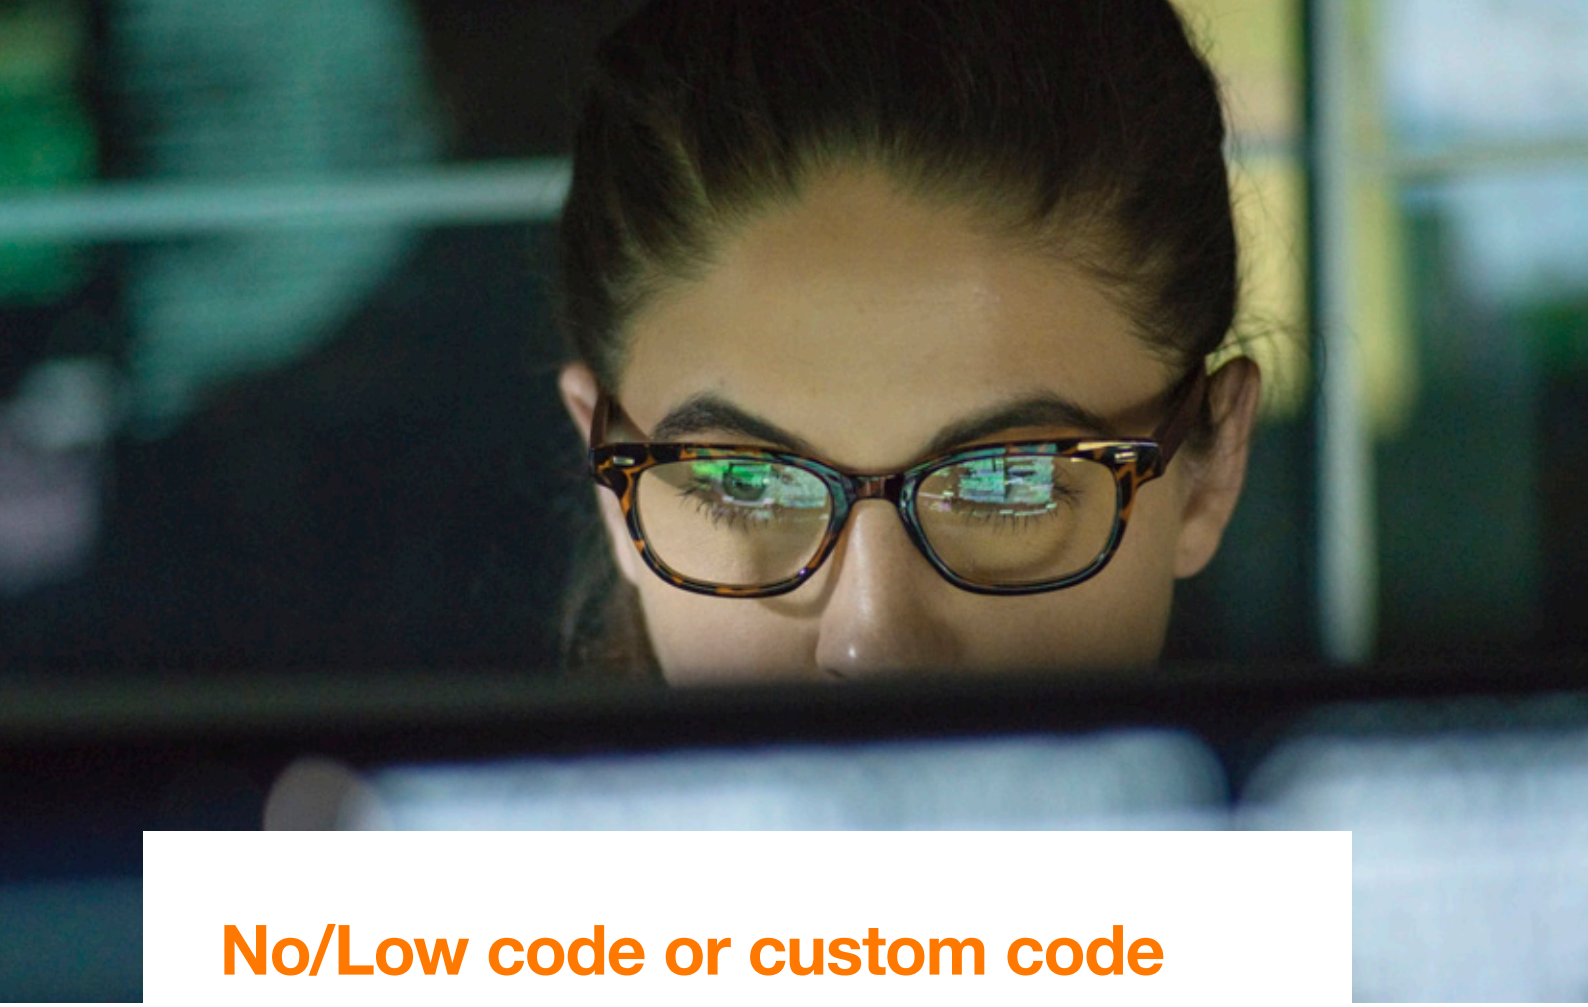

No/Low code or custom code

Custom code offers maximum control and flexibility, but it is resource-intensive to design, secure, integrate, and operate. A No/Low Code approach accelerates delivery by leveraging vendor orchestration, integrations, UI, security, and governance.

- No/Low code works best when the goal is speed for repeatable productivity scenarios (e.g., summarizing, drafting, basic analysis, simple automations) and when supported connectors align with your stack and compliance needs. The trade-off is black-box behavior and less fine-grained control.
- Choose custom code when the use case is strategically differentiating, requires specialized agents/tools, deep observability, strict data-residency/privacy controls, or nonstandard integrations with core systems.
- In practice, you can take a hybrid approach, start with no/low code to achieve quick wins, then add custom services or agents where performance, control, or compliance demands it.

Favor platforms that support exportable artifacts and model or tool portability to minimize lock-in and keep future options open. As a rule of thumb, if speed, standardization, and existing connectors are the primary drivers, and the workflow is not a source of competitive advantage, go no/low-code-first. When control, specialization, or strict compliance requirements lead, plan for custom or hybrid approach. As with orchestration, start simple and add complexity only when results or risks justify it.

No/low-code tools
(e.g., Microsoft 365 Copilot, Copilot Studio) are purpose-built for enhancing employee productivity and specific low-level processes; they become limiting as control, flexibility, or complexity rises, and tailored integrations are required.

Insights in action

Agentic AI moves AI from outputs to actions; but the winners will pair a business-first focus with engineering discipline.

Start with business outcomes.

Target a concrete process, demonstrate value, then scale.

Baseline, then right-size.

Establish performance with the most capable model; test smaller/specialized models and design for easy substitution.

Design for interoperability.

Prefer open, standards-based connections for tools/actions and agent-to-agent coordination to reduce custom glue.

Treat context and prompts as first-class assets.

Keep them simple and structured, version and test them, and manage context to prevent overgeneralization or missed signals.

Keep orchestration simple, until it isn't.

Start with single agent/workflows; add multi-agent roles only when needed. Make roles, message contracts, and feedback loops explicit.

Evaluate what matters.

Tie GenAI metrics to business KPIs and ensure continuous monitoring with HITL review for high-risk or low-confidence cases.

Choose no/low code vs.

custom code pragmatically. Favor solutions that minimize integration complexity and fit your existing stack, potentially through a hybrid approach.

As AI systems gain greater autonomy in reasoning and action, scaling them safely demands a holistic governance model that extends beyond data and model oversight to include agent governance. This ensures systems remain within defined boundaries and aligned with organizational policies, regulatory standards, and ethical principles. The next section focuses on how governance can foster trust, transparency, and compliance while enabling confident, responsible adoption across the enterprise.

Building trust: Governance for agentic systems

AI governance evolves alongside advancing AI capabilities. As generative AI progresses from producing outputs to acting autonomously, organizations are witnessing a shift from early RAG-based knowledge assistants to agentic systems capable of executing business processes.

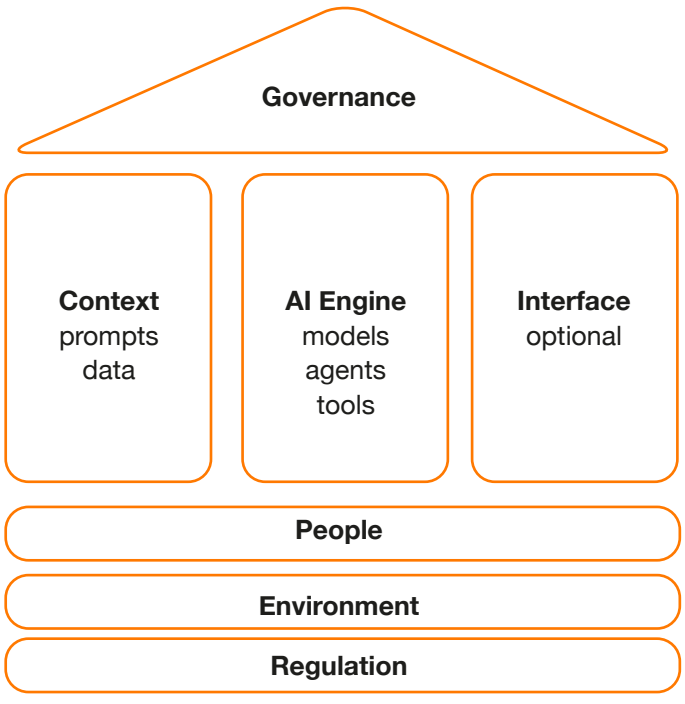

This framework illustrates how governance oversees the full ecosystem, spanning context, AI engines, and interfaces, while being grounded in people, environment, and regulation. It emphasizes that responsible governance is both structural and relational, connecting technical components with human and societal dimensions.

These extended capabilities bring opportunities, but also significant risks. With the rise of AI-related incidents, a robust governance framework becomes vital: not as a brake on innovation, but as a foundation that promotes responsible adoption, aligns AI with business strategy, and safeguards stakeholders in an environment of accelerating technological change.

AI governance defines how organizations ensure the responsible use of AI through policies, processes, accountability structures, and technical

safeguards that span the AI lifecycle, development, deployment, and maintenance. A well-designed governance model aligns operations with ethical standards and regulatory expectations, ensuring reliability and trust.

The complexity of governing agentic systems lies in their multi-layered nature: context, models, agents, tools, interfaces, infrastructure, and, crucially, the people who interact with or are impacted by them. Effective governance must consider all these components, as well as the broader regulatory and environmental landscape.

Agentic AI governance can be visualized as a layered architecture, where governance serves as the overarching structure ensuring alignment, safety, and accountability across all system components. Each layer, from the context that informs the system, to the AI engine that acts, and the interface that mediates human interaction, is influenced by and dependent on the people, environment, and regulatory frameworks surrounding it.

“With the rise of AI-related incidents, a robust governance framework becomes vital.”

Governance principles

Governance is not a constraint, it is a balance between innovation and accountability, providing a roadmap to safely integrate evolving AI capabilities into enterprise operations.

Core governance principles

1 Accountability

Human oversight must remain central. Systems should operate under clear accountability, ensuring control and traceability.

2 Transparency

Users must understand how the system functions, is deployed, monitored, and managed, and always be informed when interacting with AI.

3 Reliability & safety

Systems must be robust against misuse and errors through continuous monitoring, feedback loops, and performance metrics calibrated to acceptable thresholds.

4 Privacy & security

Protecting personal data and upholding consent and privacy rights.

5 Sustainability

AI systems should be designed and operated with awareness of their environmental footprint.

Strong governance establishes the foundation for responsible AI adoption

Trust, risk mitigation, and security must be embedded from the start by considering:

- Alignment with organizational strategy and stakeholder needs.
- Regulatory context (e.g., EU AI Act, ISO standards) and role definitions (provider, deployer, distributor).
- Risk classification under applicable frameworks (unacceptable, high, limited, minimal).
- A holistic design approach that accounts for technological, human, and environmental impacts.
- Evaluation frameworks spanning multimodal capabilities (text, image, audio, agents).
- Proactive risk management through threat modeling, audits, and red teaming.

Strong governance establishes the foundation for responsible AI adoption, but as systems evolve from static models to autonomous, context-aware agents, governance must evolve too. Traditional principles of accountability, transparency, and safety remain essential, yet agentic systems introduce new dynamics.

Governance for agentic systems

Agentic systems go beyond output generation. They plan, act, and interact with other systems. Governance, therefore, must be dynamic, adapting to evolving operational realities, tools, and contexts.

Key considerations

Purpose alignment

Each agent and tool must serve a well-defined function aligned with business objectives. Roles, permissions, and performance indicators should be explicit and continuously monitored.

Govern by design

Embed compliance, privacy, and responsible AI principles into the system architecture from inception, scaling governance with maturity, not over-engineering early prototypes.

Impact assessment

Evaluate societal, ethical, and business impacts early, integrating findings into design and deployment decisions.

Data minimization

Limit data access and processing to what is necessary, using context scoping and role-based access controls to ensure accuracy and recency.

Scope enforcement

Define and technically enforce levels of autonomy, ensuring agents operate strictly within their assigned boundaries.

Agent catalogue

Maintain visibility into agent creators, usage, performance, and alignment with KPIs to enable lifecycle management and accountability.

Continuous monitoring

Establish ongoing evaluation against metrics for performance, security, and compliance, feeding insights back into refinement cycles.

Testing

Conduct pilot trials with end-users to gather actionable feedback before full deployment.

Documentation

Capture design decisions, dependencies, and operational context comprehensively to ensure reproducibility and traceability.

Oversight framework

Compliance & risk

- Risk tiering, documentation (model and agent cards), audits, and impact assessments.
- Alignment with evolving regulatory and ethical standards.

Quality & reliability

- Retrieval and grounding accuracy, coherence, task success, and recovery from failure.
- Monitoring of latency and throughput performance.

Safety, security & privacy

- Prevention of harmful content or bias.
- Tool and code security, secret management, and privacy controls.
- Authentication, authorization, and access management.

Agentic controls

- Oversight of task adherence, intent resolution, and tool invocation.
- Guardrails including sandboxing, rate limits, and spending caps.
- Human override and escalation mechanisms.

Observability & operations

- Real-time monitoring, tracing, and auditing.
- Red teaming, evaluation suites, and incident response protocols.
- Lifecycle management and rollback mechanisms.

User experience & sustainability

- Human-in-the-loop design, explainability, and feedback integration.
- Secure, accessible, and transparent interfaces.
- Sustainable infrastructure minimizing energy and compute costs.

“Governance for agentic systems is a continuous process, not a fixed framework.”

Governance for agentic systems is a continuous process, not a fixed framework. It goes beyond control, enabling safe autonomy and responsible innovation. By blending structured oversight with adaptive mechanisms, enterprises can empower agents to act with confidence while preserving accountability and alignment with human and organizational values. In doing so, governance becomes not a constraint but an enabler of trusted and scalable AI innovation.

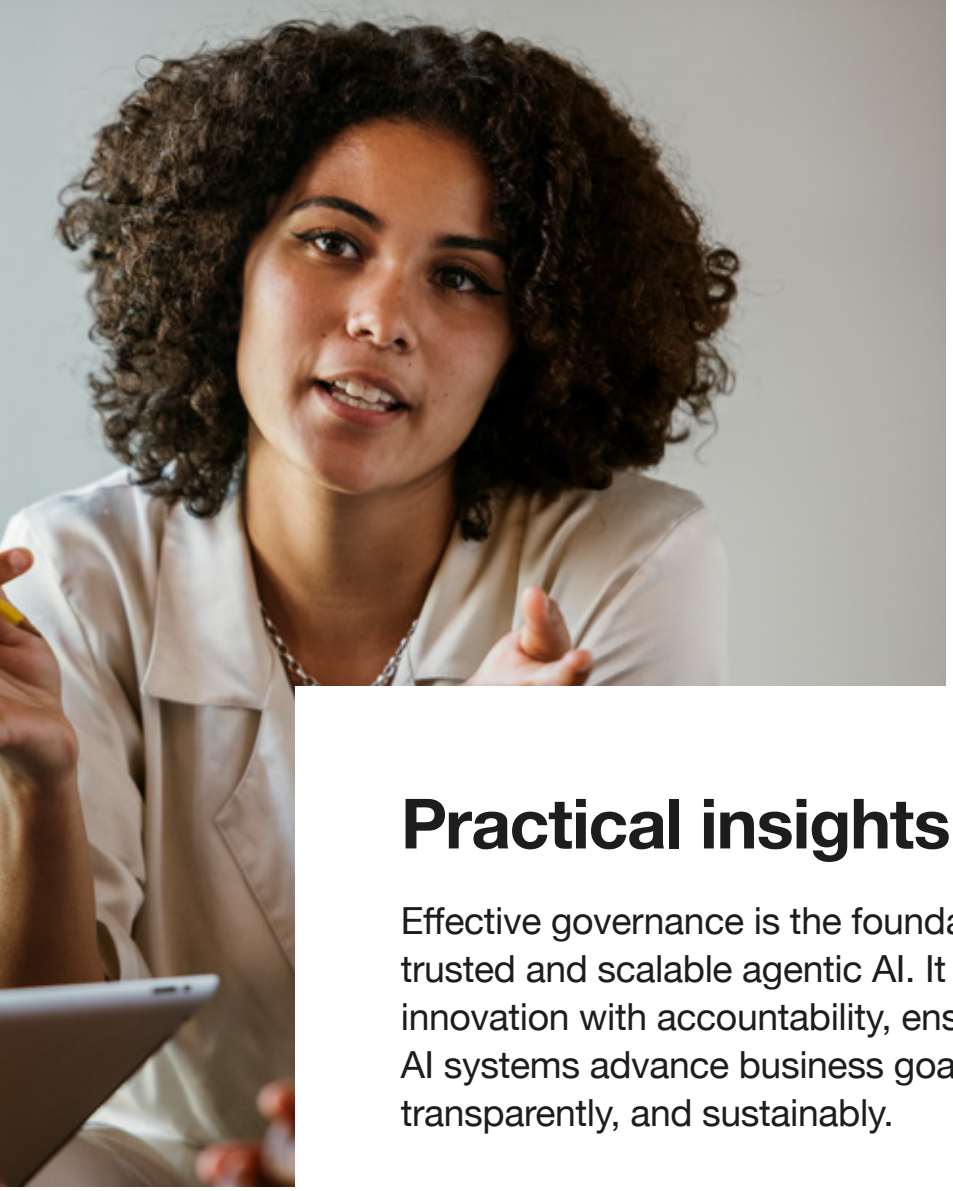

Practical insights

Effective governance is the foundation for trusted and scalable agentic AI. It aligns innovation with accountability, ensuring that AI systems advance business goals safely, transparently, and sustainably.

People first

People do not use technology they do not trust. AI must be designed, developed, and deployed around human needs and values.

Governance as an enabler

It is not an afterthought or constraint; it unlocks the potential of Agentic AI while ensuring safety, transparency, accountability, and compliance.

Human oversight remains essential

Regardless of AI progress, human-in-the-loop mechanisms are vital to maintain trust, manage risk, and uphold ethical integrity.

Inclusive creation

Today's creators include non-technical users leveraging no-code tools like Microsoft Copilot 365; governance must therefore be intuitive, embedded, and accessible.

Adaptive and continuous

Governance should evolve with AI capabilities, combining structured oversight with real-time monitoring and improvement.

Culture drives adoption

Building awareness, training, and shared responsibility fosters trust and empowers people to use agentic AI effectively and responsibly.

As governance matures, its impact ultimately depends on people, their trust, understanding, and confidence to work alongside intelligent systems. The next section focuses on how organizations can turn governance into culture by empowering employees to adopt agentic AI effectively, securely, and responsibly.

Adoption: Empowering people

Technology and policy alone cannot ensure adoption, culture does. The success of agentic AI depends on people: building trust through communities, training, and hands-on experimentation. Organizations must equip their workforce to engage responsibly and confidently with AI. This means fostering awareness, promoting transparency, and encouraging participation in AI initiatives. Ultimately, the path to trusted agentic AI lies in empowering people to use it effectively, securely, and ethically.

Engineering excellence and sound governance are essential foundations for enterprise AI, but they alone do not guarantee adoption. As AI systems become agentic, capable of autonomous reasoning and action, success increasingly depends on how people understand, trust, and use them. Adoption is not limited to advanced AI solutions; it begins with everyday productivity tools such as ChatGPT and Microsoft 365 Copilot, along with the habits people build around them. Agentic AI is reshaping work itself, transforming roles, workflows, and even organizational structures as capabilities move from innovation teams to the broader enterprise.

While policies and platforms establish structure, true adoption occurs when people are equipped, confident, and empowered to experiment safely within clear guardrails.

“True adoption occurs when people are equipped, confident, and empowered to experiment safely within clear guardrails.”

The human imperative

As enterprises advance from building and governing agentic systems to scaling their use, success ultimately depends on people. Change management becomes the strategic, adaptive capability that enables adoption, aligning people, processes, and technology. It combines visionary leadership, structured governance, and iterative learning to drive engagement and resilience across the organization.

Managing the change lifecycle for agentic AI requires deliberate planning, clear ownership, and continuous feedback. Successful adoption begins with empowering people, building enterprise and human capabilities, and fostering collaboration across business, IT, and risk functions.

Agentic AI reaches its full potential when people become confident, trusted participants in an AI-driven enterprise.

Thesis

Successful adoption of Agentic AI begins with empowering people, building enterprise and human capabilities, and fostering collaboration across business, IT, and risk functions.

Adoption principles

Adopting agentic AI requires more than deploying new tools. It demands a deliberate framework that aligns strategy, governance, and human enablement. The following principles outline how organizations can translate ambition into action by combining structured change management, practical enablement, and continuous feedback.

1 Strategic alignment

Ensure AI initiatives directly support business priorities and operational realities. Engage key stakeholders early to build shared ownership and sustained sponsorship.

2 Structured change management

Design a tailored approach combining structure and agility. Establish governance, stakeholder engagement, and communication plans anchored in measurable adoption KPIs.

3 Use case–driven enablement

Anchor adoption in tangible business scenarios. Translate AI capabilities into role-specific workflows that create visible value and mindset shifts.

4 Impact measurement and feedback

Measure usage, satisfaction, and productivity gains. Use feedback loops to refine initiatives and ensure continued alignment with business outcomes.

5 Capability building

Foster internal autonomy by developing champions, reusable content, and peer learning communities. Cultivate a long-term evolution in mindset and practice.

Focus training on both capability and responsibility

- Teach the full system: prompting, context, tools, data quality, evaluation, and responsible AI.
- Make “when to stop and ask” explicit through Human-in-the-Loop (HITL) guidance.
- Provide templates, use-case canvases, risk checklists, prompt kits, and testing plans.
- Recognize and reward early adopters and the reuse of approved patterns.
- Deliver specialized training to key individuals through a governance-aligned Center of Excellence (CoE).

Together, these principles form the foundation for effective and responsible adoption, linking strategy, structure, and capability into a unified enterprise approach. However, real success depends on how these principles are brought to life in practice:

Put people at the center

Those closest to the work best understand where agents add the most value, engage them early to build trust and accelerate impact.

Empower at all levels

Use no/low-code platforms to enable safe experimentation, reduce dependency on specialists, and democratize innovation.

Balance autonomy with structure

Pair freedom to explore with clear playbooks, targeted training, and embedded change management.

Co-create across functions

Bring business, IT, and risk together to identify use cases, define controls, and measure value collaboratively.

Adopt a right-fit build strategy

Choose between no/low-code, custom, or hybrid approaches based on maturity, capability, and evolving needs.

Treat adoption as enterprise transformation

Manage it as an organization-wide change program, not a technology rollout.

Adoption journey

Adoption thrives on community, training, and experimentation. A phased approach ensures momentum and safe scaling.

1. Assses

Evaluate current adoption maturity, identify key stakeholders, and align governance and goals. Conduct readiness workshops, communication planning, and baseline measurement.

Deliverables include

- readiness reports and stakeholder maps
- AI communication plan
- governance and security protocols
- baseline dashboards.

3. Activate

Deploy real-world use cases and empower champions to drive hands-on learning. Embed AI into daily workflows through workshops, coaching, and continuous feedback.

Deliverables include

- onboarding kits and dashboards
- progress reports and contextual learning materials
- champion certification programs

2. Design

Define the roadmap, prioritize use cases, and shape the change framework. Co-design enablement journeys with clear KPIs and ownership.

Deliverables include

- prioritized use-case lists,
- quick-win plans,
- KPI frameworks
- a mapped champion network

Key practices

- Establish safe sandboxes with clear boundaries and audit trails.
- Start with high-leverage, KPI-linked use cases.
- Favor no/low-code for speed; extend with custom components when justified.
- Track utilization, collect feedback, and iterate.

4. Scale

Sustain and expand adoption through reinforced behaviors and institutionalized expertise. Harden prototypes into production systems with robust operations.

Deliverables include

- impact and value realization report
- champion community toolkit
- success stories and training portal
- updated KPIs
- recommendations

Operationalize with rigor

- Define clear performance targets and cost controls.
- Integrate with identity, approvals, and audit systems.
- Maintain an agent catalog (owners, scope, metrics, costs).
- Embed risk management: incident response, red-teaming, model oversight.
- Tie technical signals to business outcomes: quality, adoption, efficiency, and trust.

Good practices

- Celebrate wins, document lessons learned.
- Build a CoE to scale standards, governance, and best practices.
- Maintain communities of practice and demo sessions.

Lessons from the field

Adoption efforts often stumble when ambition outpaces structure. Without clear planning, continuous monitoring, and coordinated control, even well-intentioned initiatives can lose momentum. Inconsistent governance, limited resources, and fragmented communication further compound the challenge, creating gaps between experimentation and sustainable value. Strengthening adoption requires recognizing these risks early and applying proven patterns that balance innovation with discipline.

Address challenges proactively

- **Visibility:** Establish metrics to show value and productivity impact.
- **Change management:** Build structured communication and engagement programs.
- **Resistance:** Create safe learning spaces and training opportunities.
- **Talent gaps:** Leverage partners to guide implementation and accelerate upskilling.

Anti-pattern	Do-instead
Tool first, problem later	Start with process pain and measurable outcomes.
Shadow experimentation	Provide sandbox environments with clear guardrails.
“Launch and leave”	Operate with playbooks, monitoring, and ownership.
Unnecessary complexity	Keep orchestration simple; scale with evidence.
Governance as barrier	Embed governance-by-design with streamlined paths.

Final reflections

Adoption of agentic AI is ultimately cultural. Equip people with the skills, patterns, and confidence to explore, learn, and improve. When enablement, governance, and measurement operate in harmony, agentic AI becomes a trusted capability, not just a promising technology.

Treat change management as a living system, with checkpoints, readiness reviews, and adaptive controls, to ensure safe and scalable transformation.

In the end, value is created through people. Adoption happens when teams are empowered to use agentic AI effectively, securely, and ethically, supported by clear guardrails, practical training, and a culture that rewards learning and responsible impact.

Equip people with the skills, patterns, and confidence to explore, learn, and improve.

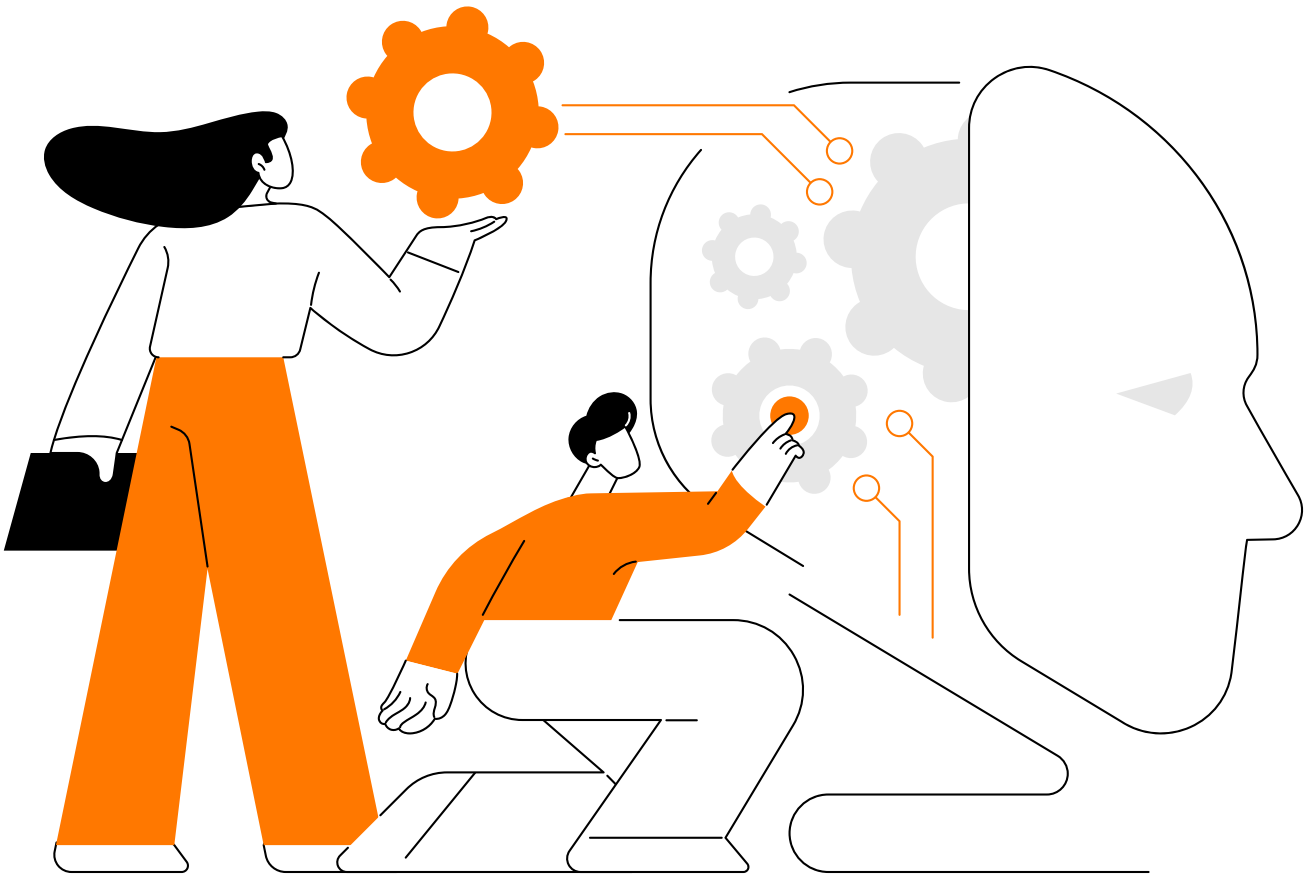

Conclusion

The evolution of agentic AI is reshaping how enterprises design, govern, and scale intelligence. Across this white paper, we have traced that progression: from defining autonomy with intent, to building scalable systems, embedding governance, and empowering adoption.

Each stage reinforces the same principle: success in agentic AI depends on clarity, purpose, architecture, and accountability. Scalable design ensures resilience and adaptability. Governance provides the structure for trust and compliance. Adoption anchors technology in people and culture. When these dimensions work in concert, agentic AI becomes more than an innovation initiative; it becomes an enduring enterprise capability.

Enterprises that choose autonomy with intent redefine the role of AI in their operations. By aligning human insight with intelligent systems, they create organizations that act with integrity, learn with context, and scale with confidence, transforming AI from a promising technology into a trusted tool.

“AI becomes more than an innovation initiative; it becomes an enduring enterprise capability.”

Contributors

Orange Business

- **Diego Olaya:** Offer Strategy Lead — Author
- **Jérémy El Aissaoui:** AI Tribe Lead — Review and Advice

Acknowledgments:

- **Renaldo Candreva:** Graphic Designer

Do you have any further questions?

Or, if you'd like to learn more about using AI to create business value from your company's data, feel free to get in touch.

Orange Business

Our joint mission is to help customers innovate and drive their business strategies in key digital domains, including Cloud, Customer Experience, Workspace, and Data & AI. We assist them on their digital journey by providing advisory, end-to-end solutions, managed services, and professional services to ensure our customers' success. We are digital natives, with innovation at the core of our business, which makes us a reliable partner close to our customers, leading them in their digital transformation challenges.

We support a wide range of industries in the private sector as well as the public sector. We have built a significant level of experience and understanding over the last 30 years in industries like Finance, Insurance, Life Sciences, Healthcare, Manufacturing, Travel & Transportation, Retail, and the Public Sector.
As always with Orange Business, our customers trust us for delivering end-to-end, sovereign, and sustainable solutions to turn their Operational Experience, Employee Experience and Customer Experience into business value.

Find out more how we can help you with your projects on:

www.orange-business.com

Follow us on

orange™ Business