



Agentic AI: The Leadership Test CXOs Can't Afford to Fail

12 Executive-level
questions that will
define your next
advantage


JOINTLY PUBLISHED BY

Spark



twimbit

Contributors



Luis Carlos Cruz Huertas

Distinguished Engineer, Artificial Intelligence and Analytics

David Chin

Chief Executive Officer, SPARK

Chang Sau Sheong

Biren Kundalia

Regional Chief Information Officer, Tokio Marine Asia

Mohan Veloo

Chief Technology Officer – Asia Pacific, China & Japan, F5

Radhey Shyam Neelakantan

Chief Operating Officer – Asia Market, Reinsurance Group of America, Incorporated

Dr Sigrid Rouam

Chief Data Officer and Head of Gen AI & Analytics & Data Management, EFG Private Bank

Narayan Venkataraman

Assistant Director Data Science and Intelligence, Changi General Hospital

Asmita Ranade

Senior Director, Affinity Analytics, Aon

Nikolay Novozhilov

Head of Data Science and Advanced Analytics, Bank of Singapore

Xin Jin

Head of Data Science & Analytics (DSA), Alexandra Hospital, Singapore

Weng Zaishan

Head, Data Analytics, MOH Holdings

Table of Contents



Executive Summary	4
Introduction	5
Twelve C-Level Questions to Prepare CXOs for Agentic AI	10
Strategic Direction	11
Business Value	18
Organisation Readiness	31
Governance and Control	45
Leading the Way	52
Conclusion	53
Contributors	54
Endnotes	55

Disclaimer:

This report is intended for informational purposes and presents insights from collaborative research and analysis with industry experts. While every effort has been made to ensure accuracy and completeness, the rapid evolution of Agentic AI means that details may change quickly. Organizations should seek expert guidance when making decisions about Agentic AI strategy and implementation. We are not liable for any actions taken based on the information in this report. The views and interpretations presented herein may evolve as the field develops. Readers are encouraged to use this information as a starting point for their own discussions and due diligence.

Executive Summary



PREPARED BY

Luis Carlos Cruz Huertas

Co-chair, SPARK Agentic AI

Advisory Council

Agentic AI marks the next major shift in enterprise technology. Systems that can perceive, reason, plan, and act independently are no longer experimental. They are entering critical business functions—and reshaping what speed, scale, and intelligence look like inside organisations.

By 2028, one-third of enterprise applications will embed agentic AI. (Gartner). 15% of operational decisions could be made autonomously. Forty percent of enterprises have already deployed agents. Another forty percent are preparing to follow. (Anthropic, Mck)

Agentic AI is moving faster than organisational structures, workflows, and governance models are prepared to change—especially as leaders are beginning to settle down with generative AI.


With this new horizon, the leadership challenge is clear: organisations that orchestrate agentic AI effectively will scale faster, operate smarter, and widen the gap against slower competitors.

But deploying agentic AI isn't just about technology. It's about making faster, smarter, and risk-calibrated decisions across the entire enterprise.

This whitepaper offers a playbook for leaders ready to move beyond pilots and build lasting advantage — structured around 12 essential questions that every CXO must answer across four critical areas:

Area	The Leadership Test
Strategic Direction	Set the course before agents set the pace
Business Value	Prove impact or lose momentum
Organisational Readiness	Build teams that win with AI, not against it
Governance and Control	Stay in charge—or lose control of the systems you deploy

Introduction



Enterprise AI stands at a defining inflection point. The generative wave sparked a surge of innovation that transformed knowledge work—augmenting content creation, research, coding, data analysis, ideation, and automation. Early adopters have realised measurable returns. Now, a second wave is cresting.

Agentic AI signals a more profound shift: the rise of systems that don't just assist decision-making—they act. It sets goals, adapts to changing environments, and drives outcomes without requiring human prompts at every step. While versions of these software systems have existed for years, the natural-language capabilities of gen AI unveil new possibilities, enabling systems that can plan their actions, use online tools to complete those tasks, collaborate with other agents and people, and learn to improve their performance.

And for the first time, organisations must design for intelligence with agency, not around it. Agentic systems traditionally have been difficult to implement, requiring laborious, rule-based programming or highly specific training of machine-learning models. Gen AI changes that. When agentic systems are built using foundation models rather than predefined rules, they have the potential to adapt to different scenarios in the same way that LLMs can respond intelligibly to prompts on which they have not been explicitly trained. Using natural language rather than code, a human could direct a gen AI-enabled agent system to accomplish a complex workflow.

With this immense potential, enterprises and leadership teams across industries are placing big bets on this technology. By 2028, agentic AI will be embedded in one-third of enterprise applications. Fifteen percent of operational decisions could be made autonomously—up from nearly zero today. Already, 40% of enterprises have agents in production, and 41% are racing to catch up. IT budgets are shifting accordingly, with spending on agentic architectures growing more than 60% year-on-year.

In the past year alone, Google, Microsoft, OpenAI, and others have invested in software libraries and frameworks to support agentic functionality. LLM-powered applications such as Microsoft Copilot, Amazon Q, and Google's upcoming Project Astra are shifting from being knowledge-based to becoming more action-based. Companies and research labs such as Adept, crewAI, and Imbue also are developing agent-based models and multiagent systems. Platforms like n8n, make.com, Zapier are enabling users to build their own agentic systems with no/low-code interfaces.

But before you decide how to deploy agents, you need to understand what they truly are.

How agentic AI works

Agentic AI represents the next evolutionary leap in artificial intelligence, transitioning beyond traditional automation and generative capabilities. While traditional AI focuses on narrow tasks and predictive outputs, and Generative AI specializes in creating content and simulations, Agentic AI systems are designed to autonomously make decisions and take actions in complex environments. An AI agent is a software entity capable of perceiving its environment and acting upon it with a high degree of autonomy. It is a system designed to reason through complex problems, interpret and create actionable plans, and execute these plans using a suite of tools.

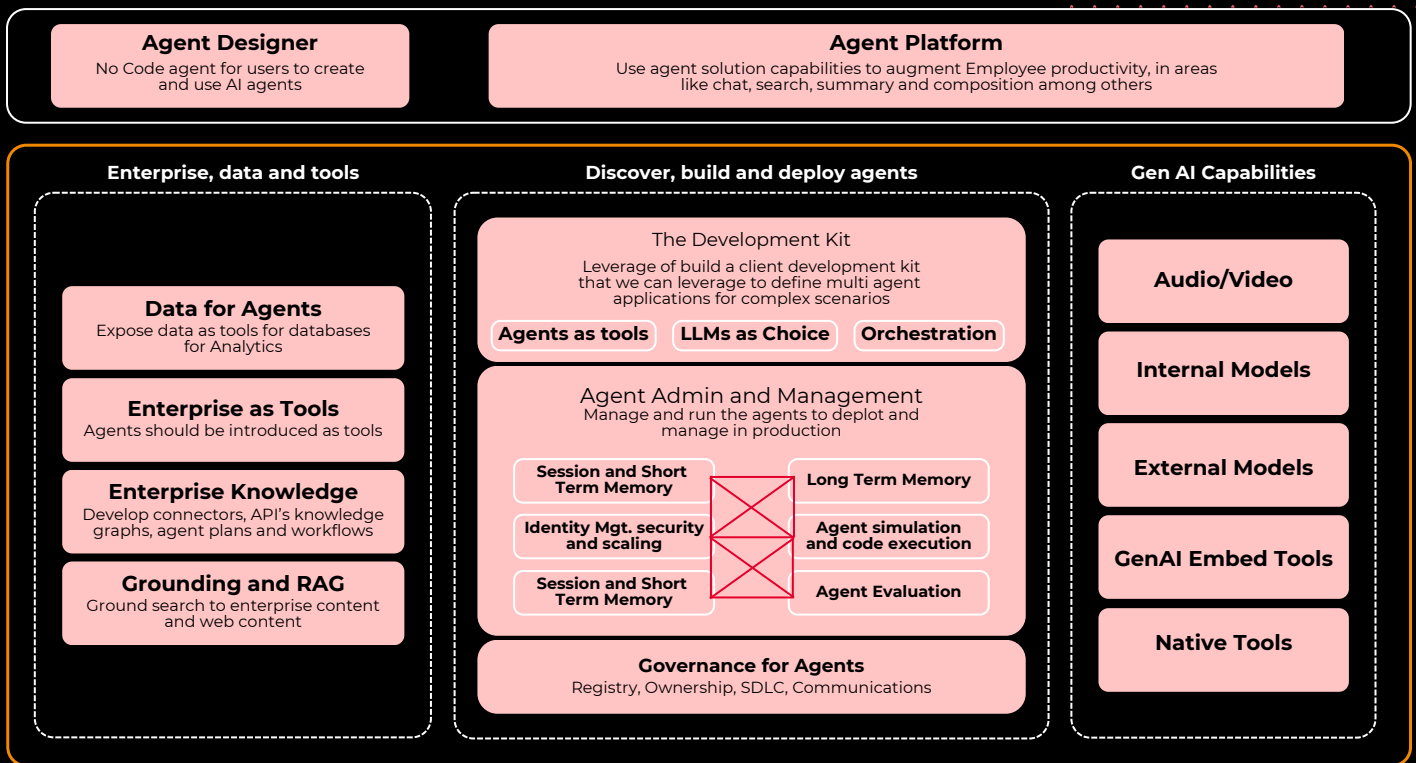
Key differences from traditional AI vs Generative AI vs Agentic AI

Aspect	Traditional AI	Generative AI	Agentic AI
Primary Function	Predict outcomes or classify data	Generate content (text, image, code, etc)	Analyse, decide and act autonomously
Autonomy Level	Low	Medium	High
Interaction Mode	Static outputs based on user queries	Dynamic content generation upon prompts	Proactive, goal-driven behaviour
Learning Mode	Pre-trained models fine-tuned over time	Fine-tuning with large datasets	Continuous learning from environment
Examples	Fraud detection, recommendation engines	Chatbots, creative writing tools	Supply chain optimisation, autonomous R&D
Human Oversight	Heavy involvement	Moderate involvement	Human-in-the-loop for critical oversight
Use Case Complexity	Narrow, specific tasks	Creative and conversational tasks	Complex, multi-step problem solving

An AI agent combines perception, reasoning, planning, action, and feedback—structured into a closed loop. Unlike traditional LLMs, agents operate through a coordinated system of sensors, reasoning engines, planning modules, and effectors. They can call APIs, manipulate tools, consult memory, and adapt plans in real time. This loop allows the system to:


- Perceive the environment using APIs, databases, or real-time user input
- Interpret and reason using memory and logic modules (often built on LLMs)
- Define a goal, generate plans, and break tasks into sub tasks
- Execute those steps by interacting with tools, systems, or other agents
- Learn and adapt by monitoring the results and adjusting strategies in real time

In enterprise environments, this is orchestrated by an agent workbench that enables businesses to deploy agents without deep technical expertise while managing access, prioritises actions, monitor performance, and enables improvement.



The Agent Workbench serves as the central ecosystem for agent lifecycle management and governance, composed of:

- **Enterprise Data and Tools Access:** Native integration with internal enterprise systems ensures agents have secure and compliant access to operational data and software.
- **Data for Agents:** A curated set of datasets, APIs, and enterprise knowledge bases optimised for agentic tasks.
- **Ecosystem Tools:** External tools and services that agents can leverage, enhancing their capabilities and interoperability.
- **DSS Development Kit:** Development and simulation tools to safely build, test, and deploy agents (DevOps for agents).
- **Code Sandbox and Simulation Environment:** A secured environment where agent code and behaviors are tested rigorously before production deployment.
- **Agent Management:** Centralized monitoring and control of all deployed agents, including lifecycle tracking, performance analytics, and health monitoring.
- **Governance for Agents:** A comprehensive governance layer that manages agent roles, permissions, policies, and ethical standards.
- **Session and Long-Term Evaluation:** Regular evaluation cycles to ensure agents remain aligned to business goals, ethical standards, and operational reliability.
- **Native Tools Usage Monitoring:** Oversight to ensure agents are only accessing and utilising enterprise tools within approved parameters.



Enterprise AI is quickly shifting from models that generate content to systems that “plan and complete goals with limited supervision.” This dynamic planning-reasoning-execution loop is what enables agentic AI to deliver business value at scale:

- Multi-agent collaboration for complex, cross-functional tasks
- Domain-specific reasoning for high-stakes decision support
- Adaptive orchestration that integrates into existing systems without automation workflows breaking easily

The challenge though is not entirely technological. It’s structural. It demands leadership choices about how the organisation is structured, governed, and incentivised to take advantage of it. The performance gap won’t be defined by who has access to the most powerful models but winners are those who knows how to design, deploy, and scale agentic systems with intent.

“Returns from AI take many forms. But you need a blueprint—clear goals, business-aligned use cases, and measurable KPIs.” says *NVIDIA’s VP of Enterprise AI, Justin Boitano*. The design of that blueprint demands a new enterprise mindset.


Leadership teams now face a web of tensions that previous transformation efforts have not prepared them for. Here are some examples:

- Delegation vs. Control: How much decision-making authority can safely move from humans to agents?
- Speed vs. Stewardship: How do you deploy Agentic AI fast enough to compete without outpacing your ability to govern it responsibly?
- Empowerment vs. Redundancy: How do you reframe human roles when AI agents can execute faster, adapt quicker, and learn on their own?
- Value Creation vs. Risk Amplification: How do you prove business impact while managing legal, ethical, and brand exposure from autonomous actions?

It’s no longer enough to ask what AI can do. Leaders must ask what AI agents should do—and what only humans must do. Most organisations today are experimenting tactically—but very few are leading strategically.

The noise surrounding Agentic AI is deafening. Vendors overpromise. Internal experiments fragment. Risk frameworks lag behind reality. Governance models designed for static systems struggle with dynamic, self-improving agents.

Without strong leadership, agents amplify weaknesses just as quickly as they scale strengths. While agentic systems can scale automation, they can also amplify data bias, policy misalignment, or hallucinated outputs. Poorly designed agents may act without traceability—or worse, make decisions misaligned with enterprise values. This makes governance, observability, and human oversight not optional, but foundational.



Leadership matters now more than ever.

Agentic AI is clearly accelerating into the enterprise mainstream. But most organisations are still in pilot mode, lacking clear frameworks to move from experimentation to systemic value. What's needed is a new strategic conversation—one that elevates agentic AI beyond hype or tooling into a C-suite agenda focused on outcomes, accountability, and advantage.

And transformation starts with asking the right questions. What does success look like in a world where AI can act with intent? Where do humans add the most value when machines can reason and execute? How do you scale capabilities while safeguarding trust?

This whitepaper is structured around essential domains to help you evaluate, shape, and lead your organisation's next wave of AI transformation—strategic direction, business value, organisation readiness and governance and control.

As enterprise adoption accelerates and the scope of agentic capabilities expands, the gap will widen between those who lead and those who follow. The organisations that succeed won't be those with the biggest models—they'll be the ones that ask better questions, act decisively, and build agentic systems designed for the long game.

Twelve C-level questions to prepare CXOs for Agentic AI



Agentic AI is not just another wave of automation. It represents a fundamental shift in how enterprises operate, compete, and innovate. For CXOs, success will depend less on technological mastery and more on leadership clarity—knowing what questions to ask before springing into action.

These twelve questions serve as a strategic compass to help C-suites navigate the risks, opportunities, and organisational changes that Agentic AI will unleash.

We organise them into four critical leadership domains:

1. Strategic Direction

- a. What strategic competitive advantage can agentic AI unlock for my organisation?
- b. How do I align agentic AI with our goals and ensure leadership sees its full value?
- c. How will skills and talent in my organisation need to change?

2. Business Value

- a. Have I defined the right outcomes and removed barriers to capture the impact of Agentic AI?
- b. Which of my domains offer the biggest opportunities beyond POCs to drive impact and build confidence?
- c. What is the true investment needed to implement and scale Agentic AI in my organisation?

3. Organisation Readiness

- a. What are the implications on my data and technology stack?
- b. How do I build the right structure and culture to instil confidence in my teams to embrace Agentic AI?
- c. Should I build or buy Agentic AI?

4. Governance and Control

- a. Have I defined the right principles and roles required to govern AI agents?
- b. What security protocols, tools, and practices are needed for my organisation to build trust towards Agentic AI?
- c. What new expectations will our stakeholders have of us when scaling AI agents?



Strategic Direction

What competitive advantage can agentic AI unlock for my organisation?

Agentic AI introduces enterprises to more than a smarter way to automate. These proactive systems introduces a new unit of execution: smarter systems that plan, reason and act independently towards goals—with minimal human guidance. (HBR)

In the past, automating something as simple as planning a business trip would involve complex multiplicity that spans different airlines, hotel rewards programs and daily schedules, all coordinated across different platforms. Today, these agents enable the automation of complex use cases with variable moving parts and kickstarts the transformation of operational processes which requires planning, decision-making, coordinating, and adapting to complex environments.

Automating what couldn't be automated before


Agentic AI enables full-process execution across domains that were previously considered too variable or too complex to automate with AI. One prominent example is pricing optimisation. Setting prices accurately often required deep understanding of supply chain costs, material availability and logistics. AI agents not only analyse historical sales data, current supply chain data such as inventory levels and transport cost to calculate price elasticity, visualise potential market response and provide recommended pricing narrative for strategy teams.

This capability unlocks the long tail of business processes—where automation was previously too expensive, brittle, or time-consuming to justify. The foreseeable result is lower cost-to-serve across more workflows, particularly in areas where interactions are highly variable such as finance, operations, and customer service.

Accelerating decision velocity

Agentic AI systems are not the dumb, passive tools teams were used to. They make decisions, iterate in real time, and coordinate across systems. In another supply chain example, agentic AI during a severe weather event can do more than just issue an alert. It could identify parts of the supply chain that may be affected, such as shipping routes, distribution centers, or critical seasonal products. The AI could then pull long-term weather forecasts and integrate this data with supplier information to assess the potential impact. Finally, it could suggest mitigation strategies, such as rerouting shipments or adjusting stock levels in certain regions, ensuring that the retailer's supply chain remains resilient and efficient.

These AI systems have greater cognitive reasoning, so they can continuously act, react, and adapt—without waiting on humans—they compress time-to-decision across functions. Agentic systems are also designed to quickly learn a company's human and brand values, ensuring that these are aligned with decisions and actions. This becomes a strategic advantage in fast-moving markets where speed is no longer a nice-to-have, but a precondition for relevance



Rewiring workforce capabilities

Agentic AI doesn't replace teams. It scales them. Instead of a human doing every task in a workflow, they supervise agents that are explicitly designed to carry out granular tasks independently. In knowledge work this could mean information retrieval, analysis, content writing, etc. This creates operational leverage and greater degree of specialisation—where one person can deliver the output of many as agents can be propped up easily.


For example in sales, Agentic AI systems could dramatically liberate sales teams from a mass of emails, paperwork, and other mundane but necessary administrative tasks. Salesforce's Agent Force interpret customer messages, recommend follow-up actions, book meetings, answer questions, and generate responses that are attuned to the company's brand voice—allowing sales agents to focus on finding and developing leads.

Enabling innovation at scale

Better understanding and judgement of content coupled with execution capabilities make Agentic AI one of the best candidates to accelerate experimentation and innovation in any industry. The most straight-forward use cases is for agents to scan and analyse both vast and niche research topics in the fraction of the time it would take humans. Platforms like Manus and deep-research capabilities in tools like ChatGPT accentuate this trend.

To take it further, humans could work together with multiple agents like in SciAgents, a multi-agent developed by a team of researchers at MIT, to identify a novel finding. The outcomes are endless as they mean developing a new insect repellent, an enhanced bio-material that combines existing materials or even new product features for a social media app trailed and tested by an army of agents. (HBR, arxiv)s

“You can build a smarter version of the same company—or you can build a different kind of company altogether.”



Strategic Direction

How do I align agentic AI with our goals and ensure leadership sees its full value?

Agentic AI will not succeed as a side project. To deliver sustained impact, it must be mapped to strategic and operational priorities—from boardroom meetings to frontline execution. That means organisation alignment around Agentic AI isn't simply a communication and change management task. It's a design requirement that starts from the top.

Start from outcomes, not capabilities

Too many AI efforts begin with the question, "What can the model do?" The better question is: What is the business trying to achieve—and what's slowing it down? Agentic initiatives gain traction when they're tied directly to business-level priorities: improving margin, accelerating cycle times, increasing NPS, growing wallet share.

Top-performing companies start by aligning AI investments with functions where pain points and potential gains are already well defined—then build system-level support around those goals. (Anthropic)

Translate vision down the stack

Enterprise priorities are often defined abstractly: "enhance customer experience," "improve operational efficiency," "scale revenue without scaling headcount." These aren't use cases—they're ambitions. The role of AI leadership is to translate these ambitions into concrete, measurable applications that business units can operationalise.

For example:

- In marketing: agents that dynamically localize content by market and customer tier
- In sales: agents that auto-summarize opportunities and trigger next-step nudges
- In operations: agents that triage incidents and coordinate resolution across systems

Each use case must reflect the language, workflows, and incentives of the business unit it touches.


Embed into workflows—not around them

Research shows that enterprise agents fail when they're bolted onto existing processes as separate tools. To align with real priorities, agentic systems must integrate into the flow of work—not as a dashboard to check, but as a co-worker to rely on. (McKinsey)

That means AI programs must work closely with process owners, line managers, and IT architects to rewire workflows in ways that augment teams without disrupting accountability and reporting lines.

Build shared ownership across functions

Alignment breaks down when AI is seen as "IT's job" or "an innovation team initiative." Agentic AI challenges traditional org boundaries—it crosses sales, ops, finance, and customer experience. To



scale successfully, ownership must be joint: business units shape the use case, technology teams build the system, and leadership defines what success looks like.

This is why several organisations are shifting toward cross-functional agent teams—small squads made up of domain experts, AI engineers, and workflow designers. Their goal: ensure each agent reflects business priorities, operational logic, and compliance requirements from day one.

Building leadership alignment on value

Many leadership teams still think of AI as a tool for efficiency—faster responses, better analytics, smarter automation. But Agentic AI represents a much bigger shift. It's about transforming how work gets done, who makes decisions, and what roles humans play. To unlock its full potential, leaders need to stop thinking of it as “just another tech initiative” and start seeing it as a new strategic operating model. Here's how to get them there.

1. Link it to what really matters to the business

Executives don't buy technology—they buy results. For leadership to care about Agentic AI, it needs to be framed not as an innovation project, but as a direct contributor to growth, margin, speed, or risk reduction. Show how it changes the economics of critical business processes, not just how it saves time. When the discussion shifts from “how it works” to “where it moves the needle,” it earns the attention of the C-suite.

2. Use real use cases to make it tangible

Agentic AI becomes believable when it's visible. Leaders don't respond to speculative benefits—they respond to working examples that solve business pain points. Well-chosen pilots grounded in measurable outcomes help leadership teams see what's possible. A small, operational pilot that shows impact will do more to shift mindsets than any strategy deck or technical demo.

3. Build trust in the technology

For leadership to back Agentic AI at scale, they must trust it. And trust starts with transparency. Leaders need to understand what the agents are doing, how they're making decisions, where the human oversight sits, and how risks are being managed. To gain confidence, they need to see that Agentic AI systems have boundaries: human-in-the-loop checkpoints, auditability, fallback mechanisms, and transparent decision paths. Governance structures that offer visibility and accountability remove uncertainty and lower perceived risk.

4. Let leaders experience it themselves

No amount of explanation replaces direct experience. When leaders interact with an agent—delegate a task, see a workflow executed, receive an output—they begin to intuitively understand the shift in capability. That hands-on exposure is often the inflection point where AI moves from abstract to actionable in the leadership mind. It demystifies the technology and opens the door to bigger ambition.

Ask the hard question: **Does my leadership team understand the full potential of agentic AI?**

They understand the hype, maybe the risk, and often the technical barriers. But few see what this technology really enables: a redefinition of how decisions are made, how work is distributed, and how organisations create value. (refer to Q1: What competitive advantage can agentic AI unlock for my organisation?)

Strategic Direction

How will skills and talent in my organisation need to change?

Agentic AI is set to reshape organisational structures, dramatically flattening traditional corporate hierarchies. It changes who does the work, how the work gets done, and what human capability looks like in an increasingly autonomous enterprise. While some roles will disappear, most will be continue to be redefined with this wave of innovation. And the organisations that get ahead will treat this not as a cost-cutting exercise. They will recognize it as a talent strategy—reinvesting savings from automation into building more adaptive, higher-leverage teams.

Running agentic systems isn't cheap. It's not unusual for advanced agent setups to cost upwards of \$10,000 per month—enough in many markets to hire multiple employees. The real question isn't whether agents are "cheaper" than people. It's whether they enable your teams to do work that's smarter, faster, and fundamentally more valuable.

From pure execution to oversight

The first shift is in functional work. As agents take on repetitive, low-value tasks—summarising documents, analysing trends, coordinating workflows—humans move up the stack. Roles shift from execution to direction, from doing the work to managing performance, escalation, and refinement at the human-in-the-loop gate.

A marketing executive might no longer pull data and build slides—but instead oversee a set of agents that generate campaign insights, flag anomalies, and draft performance reports.

This dynamic increases productivity, but it also demands new skills in delegation, validation, and interface management—skills most organisations haven't been able to institutionalise yet.

The rise of human-agent teaming


Agentic systems are not plug-and-play bots. They require context, oversight, and collaboration. This gives rise to a new class of roles focused on managing the relationship between humans and agents:

- Agent trainers: shaping behaviour through fine-tuning, RAG, and reinforcement
- Escalation architects: designing when and how agents hand off to humans
- Governance specialists: ensuring ethical guardrails and auditing decision traces
- Workflow designers: embedding agents into real-time business processes

In successful deployments, these roles don't exist as isolated technical functions—they're embedded within business units, aligned to specific outcomes. The agentic enterprise will be one where every employee acts as a manager overseeing agents.

(Adamstreetpartners)

“The future org chart includes agents. But it also includes people whose job is to make sure the agents do theirs.”



Talent pipelines must adjust at every level

Junior roles are also the first to shift. Entry-level analysts, coordinators, and assistants are being augmented—or replaced—by agentic systems that complete their tasks faster, cheaper, and with greater consistency.

This poses a long-term risk: if early-career roles disappear, how do you train future leaders who've never been exposed to foundational business processes?

To adapt, leading organizations are:

- Creating hybrid rotational programs where junior staff co-pilot with agents
- Redefining onboarding to include agent fluency and prompt-based workflows
- Designing “human escalation” tracks where judgement, not repetition, is trained


Meanwhile, the most valuable mid- and senior-level roles will increasingly require systems thinking, agent orchestration, and AI-native decision-making. Traditional business intuition will no longer be enough.

A new definition of digital fluency

Agentic AI raises the baseline for digital competencies and raises the ante for organisations competing on digital advantages. Employees must now not only understand how autonomous systems operate—what they can and can't do, how to prompt them effectively, how to audit their outputs, but more importantly, how to work with them in fast-moving environments and make the most of out of them.

This requires a reset in enterprise L&D. AI fluency can't be relegated to isolated courses or AI teams. It must become a core managerial and professional competency. Organisations that fail to navigate these uncharted waters risk more than just decreased employee satisfaction or the loss of critical human skills. They face potential reputational damage and a decline in long-term competitiveness in an AI-driven business landscape.

Top performers and early beneficiaries of enterprise Gen AI are already investing in cross-functional training, embedded experimentation, and role-specific agent support to ensure every employee can co-create with AI—not just coexist beside it. (Anthropic)



Business Value

Have I defined the right outcomes and removed barriers to capture the impact of Agentic AI?

For all the hype surrounding agentic AI, one question matters above all: Where, exactly, will the returns show up?

Enterprise leaders cannot afford to treat agentic AI as an experimental playground. They must see it for what it is — a new model for scaling productivity, compressing costs, lifting revenue, and expanding competitive capabilities.

But to capitalise, they must first understand the categories of value agentic AI unlocks — and the real-world examples already emerging.

Doing more with less

At its most immediate, agentic AI accelerates the work already being done — stripping time, repetitive effort, and human dependency from business processes.

Consider Amazon. By deploying AI agents to autonomously handle code refactoring and routine updates, Amazon freed up the equivalent of 4,500 developer-years of work annually, translating to roughly \$260 million in productivity gains. Similarly, at Power Design, an AI helpdesk agent, HelpBot, now resolves common IT requests like password resets and software provisioning without human involvement — saving the IT team over 1,000 hours of support work in its first year. (Ark Invest, Moveworks.com)


In both cases, agentic AI augment human work as broader automation takes places. These agents deployed at scale will transform various process and workflows by absorbing low-value tasks entirely while creating a better experience for employees and customers.

Cutting waste, optimising cost

Beyond raw productivity, agentic AI also optimises operational inefficiencies — driving harder cost savings.

UPS's dynamic routing system, ORION, exemplifies this shift. Instead of static delivery plans, ORION agents continuously adjust routes in real-time based on traffic, weather, and delivery volume. The payoff: 100 million fewer miles driven annually, resulting in \$300 million in direct savings and a significant reduction in carbon emissions. (CDOTimes) Google, too, leveraged agentic AI to slash operating costs — using DeepMind's systems to cut cooling energy usage by 40% in its data centers, boosting total energy efficiency by 15%. (WEForum, Google harness).

Over time, automation like this won't just lower operational costs — it will change the cost structure of entire departments. Companies that act now will be better positioned to stay profitable under pressure.



Turning speed into growth

The more ambitious payoff lies in revenue growth — where agentic AI not only cuts costs but actively fuels topline impact.

At Klarna, automation of customer service didn't just reduce expenses. It helped drive a 23% year-over-year revenue increase, by improving customer responsiveness and satisfaction. Similarly, Spirent Communications used an AI sales intelligence agent to provide real-time competitive insights to its reps, boosting sales productivity by 5% — translating to accelerated deal cycles and topline growth.

Revenue growth from agentic AI doesn't come from just selling faster — it comes from improving the buying experience and decision speed. Organisations that use AI to remove friction for customers and sales teams will not only win deals faster but capture more value over time.

Doing what wasn't possible before

Agentic AI doesn't just make existing processes faster or cheaper or more profitable. It enables outcomes that human-driven workflows physically couldn't deliver — no matter how much you scaled or staffed.

A Fortune 500 insurer deployed a suite of 78 agents and did more than just automate underwriting. It cut decision times from two weeks to three hours, opening new possibilities for real-time policy issuance — something that would have been operationally impossible with traditional teams. (Ark Invest)

These are not examples of doing the same work better. They are examples of changing what "good" looks like in an industry. Creating new competitive advantages is where agentic AI moves beyond ROI. It becomes about shifting the basis of competition — delivering products, services, and protections at a speed, precision, and resilience level that competitors cannot easily replicate.


The right metrics to measure outcomes

Investing in agentic AI is not a technology bet—it's a business bet. Yet many organisations fall into the trap of measuring input-based or vanity metrics: how many agents were built, how many employees have access, or how many prompts were issued. These metrics are easy to collect but miss the point. The real question is: Are we getting better outcomes because of AI agents?

To answer that, organisations must adopt a measurement framework that tracks value creation, performance maturity, adoption, and responsible use—across both short- and medium-term horizons.

Metrics	Short-term (3-6 months)	Long-term (>6 months)
Business Value	• Task completion rate • Time saving • Throughput increase	• Revenue contribution • Experience score uplift • Productivity / efficiency gains
Agent Performance	• Steps per task • LLM call error rate • Instruction / context adherence • Response time / latency	• Agent success rate • Tool selection accuracy / success rate • Number of human requests • Cost per interaction
User Adoption	• User engagement / adoption rate • Prompt reuse and workflow standardisation	• Agent-to-employee ratio • Time-to-productivity • Interaction quality • Agent system NPS
Governance, Risk and Trust	• Human override rate • Policy compliance rate • Data lineage • Encryption standards • User trust survey score	• Audit trail completeness • Security incident rate • Third-party audit ratings • Decision transparency / feature attribution • Bias and fairness score

(Galileo, *Understanding AI Agents blog*), (Galileo, *Mastering AI Agents*), (iaitransformation.com)



Capturing value requires effort

The examples are real. The returns are real. But the uncomfortable truth is none of this happens without a shift towards actionable, value-oriented leadership. Agentic AI is not a “deploy and forget” solution. It’s a lever that yields results if enterprises make deliberate, sometimes difficult, foundational moves.

Organisations that capture value from agentic AI do six things differently:

1. Set realistic early goals focused on productivity, not customer transformation.

Today, most agentic AI pilots deliver real value through internal improvements—optimizing operations, not transforming customer experience. Organisations must avoid the trap of measuring success by external impact too early. Productivity, speed, and internal task automation

2. Start with structured, low variance use cases where agents can perform reliably.

Agentic AI is most effective in environments where decisions follow clear rules and inputs are consistent. Early use cases should focus on simple, structured workflows—like IT support, contract approvals, or scheduling—rather than complex, judgment-based processes. Giving agents manageable environments increases their success rate.

3. Plan investments around a three-year ROI timeline, not immediate returns.

Revenue and strategic outcomes from agentic AI take time to materialise. Leaders must build business cases with a long-term view, balancing short-term operational wins with the patience required for more transformational impacts to unfold over multiple years.

4. Understand critical business processes deeply before applying agentic AI.

Superficial automation delivers superficial results. Leaders must analyse and map their core workflows end-to-end, identifying choke points, handoffs, and high-friction processes where AI-driven delegation can fundamentally reshape how value is created.

5. Redesign full workflows, not just automate isolated tasks.

Building agentic AI solutions can be agile but operationalising them inside real workflows requires heavy change management. Isolated automation, like reducing server setup from two months to two minutes, mean little if adjacent workflows—like procurement approvals—still take months. Organisations must rethink end-to-end processes to remove bottlenecks comprehensively, not just improve single steps.

6. Recognise that the biggest costs are in people, not technology.

Technology costs for agentic AI deployment are relatively low. The major expenses come from skills development, training, organisational restructuring, and managing the change at scale. Leaders must allocate major portions of agentic AI business case and budgets toward building capabilities, not just upgrading systems.

Business Value

Which of my domains offer the biggest opportunities beyond POCs to drive impact and build confidence?

Agentic AI is ready to create business value, but only if organisations are surgical in where and how they deploy it. Many organisations, however fall into common traps: scattering agentic initiatives across disconnected use cases without cumulative business impact, and selecting only high-visibility, high-risk use cases that create organisational paralysis before any real value is realised. Picking the wrong domains — overly complex, politically sensitive, poorly aligned to business needs — further risks wasted time, stalled momentum, and skepticism.

Successfully deploying agents across domains that are most likely to succeed creates proof points that build confidence, unlock ROI, and position AI as a strategic growth lever. This means domain operations are well integrated with agentic capabilities, with team members collaborating closely while also offloading tasks to independent agents with high degrees of trust and accuracy.

In a world where speed matters, the organisations that build operational AI muscle the fastest will win. To get started, picking the right domains will put you ahead in the race.

A practical scorecard to evaluate domain readiness

Criteria	What to look for	Scoring
Volume of repetitiveness	Are tasks standardised and frequently repeated? Is minimal human judgment needed?	Low = Mostly unique or creative tasks Medium = Some repetition, but high variability High = Highly repetitive, routine, high volume
Availability of structured data	Is data organised, consistent, and easily accessible? Can it support AI decision-making without major cleaning?	Low = Sparse, unstructured, fragmented data Medium = Some structured data with gaps High = Abundant, clean, structured data available
Potential business impact	Would automation materially reduce costs, improve productivity, or increase revenue? Would it meaningfully affect business outcomes or customer experience?	Low = Minor efficiency gains only Medium = Noticeable but moderate improvements High = Significant cost, productivity, or revenue impact
Ease of integration	How easily can the agent be embedded into current workflows or systems?	Low = Major rework or new infrastructure needed Medium = Moderate system/process changes High = Minimal changes, easy fit
Risk tolerance	How serious are the consequences if the agent makes mistakes? Can errors be easily caught or reversed?	Low = Very low error tolerance (compliance, safety risks) Medium = Moderate risk, manageable with oversight High = High tolerance, low consequence of errors
Readiness and alignment	Are teams open to AI adoption? Is there leadership sponsorship? Are key stakeholders aligned and tech-savvy enough to support deployment?	Low = High political resistance, poor AI understanding, no sponsorship Medium = Some buy-in, but spotty technical and political readiness High = Strong executive support, high tech fluency, minimal resistance

Scorecard in action

IT Helpdesk: Low-Risk, High-Return

Traditionally seen as a cost center, IT Helpdesk now offers one of the clearest paths for early agentic AI deployment.

- Tasks are highly repetitive, workflows are standardised, and structured data from service management platforms is readily available.
- Integration into existing systems is low-friction, minimising disruption to ongoing operations.
- Operational risks are contained: errors such as failed password resets or delayed software provisioning are easily reversible without business impact.

Critically, the organisational conditions are favourable. IT teams typically exhibit strong technical fluency, openness to automation, and clear executive sponsorship—a rare alignment that removes internal friction and accelerates time to value.

Risk Compliance: High-Stakes, High-Resistance

Risk Compliance Review, by contrast, remains a high-risk, high-resistance domain for agentic AI deployment.

- While some standardised tasks exist, much of the work demands nuanced human judgment, interpretive decision-making, and discretion that today's agentic AI cannot reliably replicate.
- Data availability is inconsistent, workflows are tightly coupled to regulatory controls, and integration would require heavy system and governance redesign.
- Error tolerance is effectively zero—mistakes carry financial, legal, and reputational exposure that far outweigh potential efficiency gains.

Organisationally, the barriers are just as high. Teams remain cautious about AI interventions in regulated environments, and leadership sponsorship can be limited due to propensity to risk.

Selecting the right domains is essential to unlocking early value from agentic AI. Structured evaluation helps organisations focus on opportunities where technical, operational, and cultural conditions align for success. By starting with strong candidates, enterprises can deliver measurable impact quickly and build the momentum needed to scale agentic AI across the business.

What is the true investment needed to implement and scale Agentic AI in my organisation?

Implementing agentic AI at enterprise scale involves far more than just buying new software or paying cloud bills. Successful adoption requires investing in both technology and organisational readiness.

In fact, AI project costs can range from as low as ~\$10,000 for a basic automation pilot to over \$10 million for large enterprise-scale deployments. (Waltorn) Many executives are learning the hard way that costs can escalate quickly – a recent IBM survey found computing costs are expected to jump 89% from 2023 to 2025, with 70% of executives citing generative AI as the cause. We foresee this will only go up with the influx of autonomous agents.

Yet many organisations miscalculate the investment required. The upfront investment is just the beginning. Enterprises that fail to account for the full lifecycle—development, deployment, integration, adoption, and governance—risk cost overruns, stalled projects, and unrealized returns.

Practical ROI estimation models

To support business case development and investment justification, organisations can use the following simplified ROI models based on practical Agentic AI deployment scenarios:

Example 1: ROI Calculation for Productivity Gains

Scenario: Deploy Agentic AI in the customer support team.

- Agents resolve 500 additional tickets per month.
- Average value per resolved ticket: \$20.
- Monthly cost of agent deployment and operations: \$5,000.

Calculation:

- Total Value Gained = $500 \times 20 = \$10,000/\text{month}$
- ROI = $((10,000 - 5,000) / 5,000) \times 100 = 100\%$ ROI within the first month.

Example 2: Cost Savings from Automation

Scenario: Agents automate 3,000 hours of employee work per year.

- Average employee cost: \$50/hour.
- Annual cost of agent development and maintenance: \$120,000.

Calculation:

- Annual Savings = $3,000 \times 50 = \$150,000$
- Net Benefit = $150,000 - 120,000 = \$30,000$ net savings in the first year.

Example 3: Revenue Uplift Through New Services

Scenario: Agents enable launch of personalised digital advisory services.

- New customer acquisition: 1,000 customers/year.
- Average revenue per customer: \$200/year.
- Incremental operating cost: \$50,000/year.

Calculation:

- Gross New Revenue = $1,000 \times 200 = \$200,000/\text{year}$
- Net Revenue Uplift = $200,000 - 50,000 = \$150,000$ net new revenue annually.

To make informed investment decisions, leaders must understand not just where the money goes, but why each category of cost matters—and how they evolve over time.

Rethinking costs of Agentic AI

Certain agentic AI tools today are priced to match that of human labor. Call center agents powered by AI, for instance, are marketed at rates as low as \$3/hour—below minimum wage. This may appear attractive at first glance. In the short term, especially for task automation, this pricing model helps buyers compare cost savings directly against headcount reduction. (LinkedIn, Appendz)

But this framing is inherently flawed. AI is not labor. Over time, its pricing will not be set by human equivalence but by marginal cost: software development amortised over scale, and hosting fees that approach zero per transaction. This is the trajectory most SaaS platforms followed, and agentic AI will likely converge in the same direction.

In reality, the true cost of enterprise AI lies not in the “per hour” pricing, but in everything around it: architecture, compliance, talent, integration, and change management.

Direct cost – What Shows Up on the Budget

1. AI Platform and Licensing Fees

Many enterprise-ready AI services follow a usage-based pricing model. For example, Microsoft Copilot is priced at ~\$30/user/month, or per interaction at fractions of a cent. Salesforce's Agentforce platform charges ~\$2 per conversation. These micro-costs scale rapidly with organisational adoption, and often catch CIOs off guard. (Siroccogroup)

For organisations with high-volume use cases—customer service, sales enablement, or HR support—the aggregate cost of tokens, API calls, and usage-based licenses can quickly eclipse expectations. (AI Infrastructure) Pricing today may be competitive, but as usage grows, so do cumulative charges—particularly for firms deploying AI across thousands of employees.

Cost Signal: Over-usage of "per conversation" or "per token" services in production environments can double or triple initial cost estimates within months.

2. Cloud Infrastructure and Compute

Compute is one of the largest and most volatile cost categories. Running agentic AI on premises—especially those with real-time reasoning, multi-step task execution, or large language model backends—demands significant cloud resources. (IBM, AI Economics)

GPU or TPU instances on cloud platforms (AWS, Azure, GCP) offer flexibility, but pricing varies with workload and availability. Companies often encounter unpredictable spikes in usage, especially when models are deployed across multiple departments. For organisations with steady, high-throughput workloads, some shift to on-premise hardware—a high capex move with lower variable cost, but significant operational overhead. (Waltturn)

Cost Signal: Enterprises running LLM-based agents across product, support, and ops functions often report compute costs as a top-three AI expense category.


3. Data Acquisition and Preparation

The quality, accessibility, and structure of enterprise data directly affect AI outcomes. But getting data ready for agentic use is neither fast nor cheap. Many organisations must invest in:

- Data labelling and annotation
- Cleansing and deduplication
- Integrating fragmented systems and legacy data stores
- Establishing secure pipelines for real-time access

Storage and compliance add further layers. Cloud storage costs rise with volume and frequency of access. In sectors like finance and healthcare, costs to meet regulations (e.g., GDPR, HIPAA) are substantial.

Cost Signal: AI initiatives, not just limited to agentic use cases, often incur more cost in making data usable than in training the model itself.



4. Specialised Talent

Hiring AI engineers, MLOps experts, and domain specialists remains expensive and the biggest chunk of many business cases for agentic AI. Salaries for experienced talent often range from \$150,000 to \$300,000 annually in mature markets. (Waltorn) Even if off-the-shelf models are used, enterprises need in-house capability to tune, manage, and align these tools to business logic.

Additionally, companies must invest in upskilling internal staff—sales reps, HR teams, analysts—so they can confidently use and manage AI tools. Upskilling, while less visible than hiring, is essential for adoption and long-term scalability.

Cost Signal: Talent costs are ongoing—not just during development, but for governance, retraining, and support.

5. Model Development, Testing, and Iteration

Building or customising AI agents involves significant iteration cycles. Each version requires compute time, developer hours, and testing frameworks. The complexity compounds with every new use case added. Enterprises building proprietary models or tuning open-source agents for internal logic must also budget for tooling—MLOps pipelines, tracking systems, and deployment environments.

Cost Signal: High experimentation cost arises not from one training run, but from hundreds of refinements during tuning.

6. Integration with Enterprise Systems

AI agents need to operate inside business workflows. This requires integration with CRMs, ERPs, service desks, and messaging platforms. In many cases, legacy systems are ill-suited for AI orchestration, forcing upgrades or middleware development.

Integration is one of the most underbudgeted cost categories—often underestimated during pilot phases and only fully recognised during scaling.

Cost Signal: Integration costs can exceed model development costs when agentic AI systems need to touch multiple legacy systems.

7. Security and Audit Controls

Implementing AI agents expand the attack surface. Enterprises and security leaders must implement robust security layers: encrypted APIs, access controls, audit trails, and model risk monitoring. Especially in regulated industries, the cost of securing agentic AI systems—against data leakage, adversarial prompts, or misuse—is non-negotiable.

Cost Signal: Expect security infrastructure and audit tooling to scale in parallel with AI deployments.

Indirect cost: Hidden, but crucial for success

1. Change Management and Organisational Alignment

AI agents doesn't just make tools work—it changes how decisions are made. Deploying these autonomous systems at scale requires orchestrating new roles, workflows, and expectations. The cost of managing this transition—via training, stakeholder engagement, and behaviour change programs—can rival the cost of building the agent itself. (McKinsey, potential of gen ai in insurance)

Some leading organisations dedicate up to 50% of total investment to change management, recognising that adoption determines return.

Cost Signal: Without change budgets, adoption stalls—even when the tech works.

2. Governance, Compliance, and Ethical Risk Management


Responsible AI governance is now a board-level concern, especially as agents take on more responsibility. Enterprises must develop clear policies for accountability, explainability, bias mitigation, and fail-safes. This involves setting up governance boards, conducting audits, maintaining logs, and engaging legal counsel—all of which carry recurring cost.

Cost Signal: Governance costs are not front-loaded—they rise over time as scale and risk grow.

3. Retraining and Performance Drift Management

Agentic AI is not static. As data shifts, behaviors evolve, or business logic changes, models will require retraining. This includes engineering time, compute resources, and QA. Enterprises with multiple models must also build monitoring systems to detect drift or degradation.

Cost Signal: Post-launch maintenance can often exceed initial build costs within 12–18 months.



The total cost can also vary significantly based on five key factors:

Variable	Impact on Cost
Scale of Deployment	Larger orgs incur higher infrastructure, training, and integration costs
Industry Regulations	Healthcare, finance, and government require extensive compliance and oversight
Build vs Buy Strategy	In-house development demands talent and R&D; vendor tools incur license fees (excluding vendor lock-in)
Infrastructure Model	Cloud = variable OPEX; on-prem = fixed CAPEX + operations overhead
Complexity of Use Cases	Multistep agents with reasoning logic cost more to build, maintain, and explain

“Agentic AI costs will not stay linear — it will curve upward with scale, complexity, and dependency.” (CIO.com)

Organisations that plan for this curve — investing early across both direct and indirect cost drivers — will lead the next wave of enterprise AI transformation. Those that underestimate it will be forced into expensive, reactive rebuilds.

Organisation Readiness

What are the implications on my data and technology stack?

Deploying agentic AI presents one of the most demanding tests of legacy enterprise systems. Those who have started this journey with Gen AI would discover that these environments, often built over decades, are riddled with fragmented data, siloed tools and applications, hard-coded logic, and outdated integration patterns.

Making agentic AI work in such settings isn't about building another software tool or calling an API. It demands rethinking how data flows, how IT systems talk, and how agents system will be governed and integrated into enterprise workflows without disrupting existing operations.

Here are several key components introduced by agentic AI:

Establishing data foundations

As with all forms of AI, the important of rich, high quality data cannot be understated. The phrase "Garbage in, garbage out" sums up the realisation for many legacy enterprises when starting out generative-based projects.

Similarly, Agentic AI works best with unified access to structured and unstructured data that provide the right context for autonomous decisions. Before deploying these agents, key components need to be in place for agents to "understand" the context it operates in.

- **Data Lakes and Lakehouses:** Nearly 90% of corporate data is unstructured (Salesforce, Vector databases). Data lakehouses combine the flexibility of a lake with data management features of a warehouse to ingest and store raw structured and unstructured data at scale (documents, logs, images, etc.). This is useful for agents to output accurate decisions as it needs constant access to available data streams and learn.
- **Feature Stores:** Feature stores play a critical role in agentic AI by enabling agents to access consistent, reusable engineered variables across models and teams. They ensure alignment between training and inference data, improving agent decision reliability. (Medium, Feature stores) These can be layered over existing pipelines to standardise feature access, reduce redundancy, and accelerate time-to-production for agent capabilities without rearchitecting legacy systems.
- **Knowledge Graphs:** These overlay structure and meaning onto fragmented enterprise data, turning them into contextually meaningful knowledge. (CIO.com, knowledge graphs) By mapping relationships (e.g. customer -> transaction -> product -> region), they enable AI agents to reason and retrieve information in context. For example, a knowledge graph can link a supply chain disruption to its likely causes and stakeholders, giving an AI agent rich context to decide next steps (Enterprisesight.com) The graph can be queried as part of the agent's reasoning (often via graph-based retrieval augmented generation, or "Graph-RAG"). In legacy environments, knowledge graphs can unify disparate datasets without requiring an exhaustive centralisation exercise.

- 
- **Data Streams:** These AI agents are proactive and needs real-time data to reach to events and trigger actions. Data streams (from IoT sensors, transaction logs, etc.) can feed into the data lakes or streaming analytics platforms like Apache Kafka to capture events that agents subscribe to and monitor for state changes.
 - **Vector Databases:** Popularised by the influx of Gen AI tools, these specialised data store contains high-dimensional embeddings (vectors representation of text) that allow agents to perform similarity searches on enterprise knowledge (documents, emails) to find relevant information by meaning, not just keywords. (Salesforce, Vector database) This improves the accuracy of the agents' understanding through RAG. For example, an agent could query a vector store to find all past incident reports similar to a current IT ticket before curating a response and taking corrective action.

Orchestrating agents across systems

For agentic AI to move from prototype to production, integration is non-negotiable. These agents are not isolated models—they must take action, retrieve data, and coordinate services across the enterprise. This demands a deliberate orchestration layer that connects legacy and modern systems in real time.

Agentic AI must operate in context. That means it must "read" enterprise signals, access relevant data, and trigger downstream systems. Whether in customer service, logistics, or finance, integration determines how useful and scalable the agent becomes.

- **APIs and Microservices:** APIs are the essential access points that enable agents to interact with existing enterprise applications. In modern systems, a microservices-based architecture captures business logic in modular units—perfect for agents. For example, an AI agent resolving a customer complaint might call APIs of an online ordering portal to fetch order history, issue refunds, or trigger follow-up actions. In environments where APIs are limited, enterprises can use wrappers or gateways (e.g., SOAP-to-REST converters) to bridge legacy endpoints. Agents don't just need access—they need clean, well-documented access to invoke actions reliably. (IBM, What is Agentic AI)
- **Message Buses and Event-driven Architecture:** Workflows and processes in enterprise functions are often long-running and asynchronous. For these autonomous systems to respond quickly to changes, they need to "listen" to business triggers across applications and responding in real-time with appropriate actions. (Workato, What is Agentic Orchestration) For example, an AI agent might listen for an inventory low-stock event and then autonomously place an order with a supplier via API. An event bus or pubsub system (Kafka, JMS) decouples event producers from event-driven consumption, which improve scalability as agents only wake on events.
- **Workflow Orchestration Engines:** While agents can reason and plan, some enterprise processes are too complex or regulated to leave entirely to agentic discretion. Workflow engines (e.g., Camunda, Airflow, AWS Step Functions) define business logic in structured


- flows, enabling agents to function reliably. In an agentic architecture, one approach is a “conductor” agent (often an LLM) that supervises other task-specific agents. (IBM, What is Agentic AI) In high-stakes domains (e.g., compliance, finance), this type of orchestration guards against runaway automation by enforcing policy gates that retain visibility and control.
- **Integration Platforms (iPaaS):** Integration platforms like MuleSoft, Workato, and Boomi help connect systems without custom code. They’re increasingly built to support AI use cases—offering pre-built connectors, event-based triggers, and support for LLMs. For enterprises with existing iPaaS setups, this is a fast way to let agents interact with SaaS apps, databases, and legacy systems. These platforms also handle retries, logging, and data transformation—key for running agents at scale. Rather than rebuilding integrations from scratch, teams can extend what’s already working. (Workato)

Enabling actions with tooling and platforms

Agentic AI is more than an LLM. It’s a system that perceives, reasons, and acts. To support this, enterprises need an AI stack that goes beyond model inference. The right tools transform raw models into agents that operate reliably, safely, and at scale.

This AI/ML toolchain sits between foundational models and enterprise outcomes. It ensures agents have the infrastructure, memory, and reasoning scaffolds to be effective.

- **Model Serving Infrastructure:** After training or picking a model, for it power agents, it must be hosted on scalable infrastructure that supports high-throughput, low-latency inference. For example, a customer service agent might make thousands of inference calls per hour—requiring load balancing, autoscaling, and GPU integration. Enterprises can consider a cloud-based service (e.g. Azure OpenAI, AWS SageMaker endpoints) or on-premise deployment using frameworks like TensorFlow Serving, TorchServe, or Ray Serve. Ideally, it should also expose APIs that the agent orchestration layer calls. Organisations with existing ML platforms can integrate agentic systems by containerising new models and deploy via their existing CI/CD pipeline.
- **Prompt Management System:** Enabled by LLMs as the core module, AI agents rely heavily on prompts, sometimes complex prompt chains that take advantage of its Chain-of-Thought (CoT) reasoning capability to achieve outcomes. In enterprise AI, these prompts encode key business logic and should be treated as first-class artifacts—versioned, tested and monitored. PromptOps tools (e.g., PromptLayer, PromptHub) allow teams to store prompt templates, run A/B tests, and analyse prompt drift. For instance, if an agent’s refund prompt stops generating compliant responses after a model update, the prompt manager flags it. Learning from Gen AI projects, organisations deploying agentic AI should tightly control versioning of prompt templates, maintain a library of tested prompts, monitor prompt outcomes repetitively and create “parameterised” prompts where business data can be injected safely. (Nvidia, Mastering LLM Techniques)

- 
- **Agentic Frameworks:** Building agentic systems often involve using frameworks and protocols that provide abstractions for multi-step reasoning, tool usage and memory. Libraries such as Langchain, CrewAI, AutoGen, Haystack or newer protocols like Anthropic's Model Context Protocol (MCP) or Google's Agent2Agent Protocol (A2A) simplify development of agents by offering components for connecting to tools, managing conversational context, or implementing multi-agent planning and execution. Enterprises may adopt such frameworks to speed up development of agent behaviors, for example, choosing a framework that allows the agent to call already installed plugins or functions safely.
 - **Model Improvement Pipelines:** The models (custom or fine-tuned) powering agents will need periodic retraining and fine-tuning due to drifts in data, discovery of new edge cases or explicit user feedback. Enterprises deploying agents at scale should consider implementing pipelines that collect data from interactions logs, identify failure patterns, trigger retraining and testing. This can also include a model registry to store versions of models used. Continuous training alongside CI/CD pipeline is ideal for agentic systems that learn over time, for instance, an AI customer service agent could be fine-tuned monthly on the latest solved cases to improve its accuracy on serving complaints about current products.
 - **Observability and Feedback Tools:** Measuring the performance (accuracy, error rates) of AI agents is key to establish trust. Visibility is therefore critical. Agentic systems operate probabilistically—they won't always fail the same way twice. Enterprises need observability platforms log model calls, analyze response and decision quality, and alert on anomalies. Many MLOps platforms (Fiddler, Arize) today provide LLM monitoring features such as outlier and bias detection that prevents enterprise operating blind.

This is not the full picture. The implications of agentic AI on an enterprise's data and technology stack go well beyond data infrastructure, orchestration, and tooling—namely security, access controls, governance and policy management). But we hope this has provided clarity on where to start—and what foundational shifts are required.

Organisation Readiness

How do I build the right structure and culture to instil confidence in my teams to embrace Agentic AI?

Agentic AI shifts how work is performed, how decisions are made, and how teams interact with both machines and each other. To scale Agentic AI responsibly and effectively, organisations must evolve their structure, processes, and culture. This isn't about spinning up a task force in a corner. It's about embedding AI deeply—everywhere.

To provide a clear model for transformation, we opt for a mental model broken into three parts: organisation structure, process design and culture management.

Reimagining the org chart

Traditional org charts are built around siloed domains of expertise— finance teams handle money, marketing teams handle messaging, and so on. But agentic AI changes this equation. When expertise no longer resides exclusively in humans but is distributed across autonomous agents, leaders should start to question whether their organisation can be structured differently. (Time.com)

Instead of building around siloed areas of disciplines, leading thinkers propose designing organisations around “jobs to be done” where humans and AI agents solve for specific outcomes. (Time.com) This opens the door to entirely new structural archetypes. A provocative alternative is the Hollywood model. On a film set, the best talent — director, cinematographer, actors, editors — come together for a single project. They collaborate intensely for months, then disband. There's no fixed structure, only purpose-built teams optimised for the task at hand.

Applied to enterprises, this suggests a future where fluid, cross-functional teams — composed of humans and AI agents — form around specific business objectives, then dissolve or reconfigure as needed.

Another emerging idea is to design orgs around the strengths of agents versus humans. Leaders should ask: What tasks are best suited to autonomous agents? What requires human judgment, empathy, or creativity? In this model, organisations split functions into agent-first and human-first categories — and then assign ownership, oversight, and structure accordingly. For example, an agent might handle automated compliance checks, while humans focus on exception handling and ethical decisions.



As the notion of digital coworkers spreads, Microsoft foresee the rise of the “agent boss” — a new type of leader who manages a portfolio of AI agents like a CEO runs a startup. Every employee, in this vision, becomes a manager of agents — creating, deploying, and delegating to their own AI coworkers. It's a radical shift in mindset, requiring new capabilities in prompt design, task orchestration, and agent supervision. (CIO.com, Agents become teammates)

While future org structures remain fluid due to the changing nature of work, organisations need structural foundations today. Best practices are emerging:


- Establishing a Center of Excellence (CoE) helps organisations align high-impact agentic AI initiatives with business goals by providing shared infrastructure, centralising governance, standardising tools, data and model practices, and enabling faster, more consistent deployment of AI agents across functions, while enabling knowledge transfers across teams. (AI Multiple Research)
- Most enterprises are adopting a hybrid AI operating model: a central governance team (the “hub”) defines standards, policies, and shared infrastructure, while embedded AI “pods” within business units (the “spokes”) lead execution. (Pupila.ai) These pods are typically cross-functional — combining data scientists, engineers, domain experts, product managers, and designers — and are responsible for taking AI use cases from discovery to deployment. Each pod should have a clearly appointed AI lead who ensures alignment with business goals and acts as a bridge to the central AI Center of Excellence.

Two types of operating models are emerging:

- o Vertical AI Agent Model, where a central “Master AI” unit oversees departmental AI sub-agents, ensuring consistency, compliance, and auditability.
- o Embedded AI Pods Model, where smaller AI pods are distributed within functions (e.g. marketing, finance), reporting to local leadership and optimising for speed and business alignment. (linkedin.com, AI Agents are changing org chart)



- As part of organisational redesign, many companies are formalising AI leadership at the top. Appointing a Chief AI Officer (CAIO) — or expanding the Chief Data Officer’s mandate — signals AI’s strategic importance and ensures dedicated oversight of the roadmap, investment, and governance. In parallel, leading firms are establishing AI governance boards or ethics committees comprising stakeholders from legal, risk, and business units. These bodies define enterprise-wide policies to ensure responsible AI use and embed ethical safeguards into the fabric of decision-making.



Redesigning processes from the ground-up

As agentic capabilities grow, business workflows need to shift beyond static, task-oriented processes. Enterprises should rethink how tasks flow between humans and agents create new roles to support AI-augmented processes, and adjust decision-making protocols.

Leaders should begin by auditing key business workflows — such as customer onboarding, IT ticketing, or loan processing — to identify where agents can perform tasks and where human judgment remains essential. The goal isn't substitution, but recomposition: designing workflows where AI and humans operate in tandem, each focused on tasks they do best.

For example, an analyst agent might gather and analyze data, hand off a summary to a human manager for review, and then resume execution based on the decision. These handoffs ensure that each actor (AI or human) is performing tasks suited to their strengths — routine and data-heavy tasks to AI, and complex or sensitive judgments to humans. (rossdawson.com)

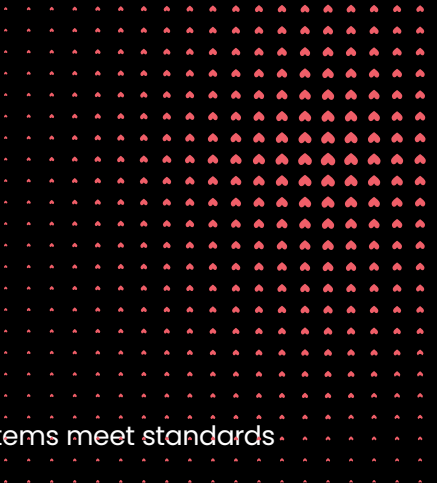
To implement this kind of “AI doing, human supervising” structure (Deloitte, Agentic AI), organisations should update process documentation, SOPs, and training. It may also involve adding checkpoints where AI outputs are audited, as well as feedback loops for AI learning (e.g. when humans override an AI decision, feed that insight back to improve the model). Over time, as confidence in AI grows, some approval steps might be relaxed (e.g. straight-through processing by AI for low-risk cases).

Two emerging patterns of human-AI teaming models are “Human-AI sandwich” and “AI with human-in-the-loop” workflows. In a human-AI sandwich, a human initiates and frames the task, an AI agent performs the execution or analysis, and then a human reviews and refines the output. The AI with humans-in-the-loop model flips the order: the AI agent autonomously handles a task or decision up to a point, and a human intervenes only at pre-defined checkpoints for approval or exception handling. (rossdawson.com)

Some organisations are even experimenting with dynamic orchestration, where an AI orchestrator assigns tasks across agents and humans in real time, depending on capabilities. To support this, new oversight roles — such as an AI controller who monitors agent activity — are becoming necessary to maintain transparency and accountability. (BenjaminLNS)

As business processes evolve with agentic AI, so too must the roles and skills that support them. Rather than replacing employees, new roles are emerging at pace:

- **Prompt Engineers** specialise in crafting and refining inputs that guide generative AI agents to deliver high-quality, relevant outputs. Their ability to frame tasks effectively can significantly influence performance.
- **AI Trainers (or Coaches)** curate training data, label outcomes, and provide ongoing feedback to fine-tune model behaviour. In an agentic AI context, they help agents continuously learn domain-specific knowledge and correct errors.

- 
- **AI Risk Managers / Ethicists** play a vital governance role, ensuring AI systems meet standards for fairness, transparency, and compliance.
 - **AI Product Managers** serve as the connective tissue between business needs and technical delivery, shaping use cases, aligning priorities, and managing outcomes.

In effect, humans take on roles of validators, explainers, and improvers of AI agents and organisations should proactively redefine job roles to emphasise how employees will work alongside AI, not be replaced by it. For instance, a financial analyst's role may be re-scoped from manually gathering data to interpreting AI-generated analytics and providing strategic insights. Customer service reps might become "AI-augmented advisors," spending less time on routine queries (as AI chatbots handle those) and more on high-touch customer interactions or supervising the chatbot's quality.

This is already playing out in the field. A large U.S. accounting firm reorganised its customer support by introducing a dedicated AI pod. AI agents now handle 65% of routine status requests — such as filing updates — while human staff concentrate on complex client needs. The firm created new roles like an AI Support Trainer, responsible for keeping agents updated with regulatory changes, and restructured decision rights: the AI handles tasks autonomously up to a confidence threshold, escalating others to humans. The result? Faster client responses, improved employee satisfaction, and better resource allocation. (CIO.com, How IT leaders)


These changes underscore a broader trend: employees must develop fusion skills — the ability to combine human judgment with AI capabilities. And the demand is clear. 94% of employees believe they can build the skills their organisations need for AI. (ciodive.com, talent upskilling) Leaders must meet this optimism with structured upskilling, clear role progression, and inclusive adoption strategies.

Develop cultural and mindset readiness

Successful adoption depends as much on cultural transformation as it does on technical deployment. Organisations must cultivate an environment of trust, experimentation, continuous learning, and ethical responsibility.

Continuous learning

In response, companies are encouraging employees to explore AI tools, attend industry forums, and share insights across teams while others are forming AI guilds or communities of practice — spaces where employees from different departments collaborate on side projects, discuss new research, and explore how AI might solve real business challenges. Others are finding that certifications and recognition play an important role in encouraging adoption. Acknowledging those who experiment with AI tools or bring new capabilities into their teams helps signal that AI fluency is valued — not just by technical teams, but across the organisation.



Crucially, AI literacy must extend to leadership. Tomorrow's managers need to make decisions with a working understanding of agentic AI's capabilities and limits. Embed these topics into leadership development programs and create mutual mentorships where technical experts and business leaders exchange knowledge — one gaining business context, the other building AI fluency.

Some companies are going further with rotational programs that place AI PhDs into marketing, finance, and operations teams. These rotations deepen domain expertise and create embedded AI champions who act as internal liaisons between technical and business units.

Across all of this, cross-pollination of skills is emerging as a theme. Software engineers are increasingly encouraged to build agentic AI fluency, while data scientists are increasingly hired not just for technical depth, but also for their ability to act as “analytics translators” — connecting insights to impact. (Ciodive.com, Employees want)

Lead by example

Leadership teams are the silent head of culture. When the CEO and C-suite openly prioritise AI agents, it sends a strong signal. For example, Shopify's CEO Tobi Lütke told employees that “using AI effectively is now a fundamental expectation of everyone” at the company (cio.com, How IT leaders) — a bold top-down mandate that set the tone for an AI-first culture. A strong leadership stance also helps overcome inertia or fear in middle management.

Building trust

Trust is the cornerstone of both employee of AI, especially as they are increasingly expected to collaborate with agents. Leaders must demonstrate honesty, competence, and care when rolling out AI, to allay fears of hidden agendas or neglect of ethics. (HBR, Employee's don't) Culturally, companies should emphasise transparency in how AI tools work and how decisions are made. This can involve explaining (in accessible terms) the role of AI in certain processes, sharing results of AI audits, and maintaining openness about AI failures or errors and what is being done to fix them. For example, if a bank rolls out an AI loan approval agent, explaining to loan officers how certain variables are analysed and has been tested for bias will make them more comfortable using it (and providing feedback).

Encourage failing

Leaders should strive to remove the fear factor that often accompanies new tech. By encouraging agile experimentation, such as hackathons, innovation labs, or “AI days,” employees get hands-on experience with agentic AI and can see its potential in a low-risk environment. Culturally, it should be clear that not every agent pilot will succeed, and that's okay — what matters is learning and moving forward. If an AI project doesn't meet its targets, teams should be congratulated for what was learned, and those insights should be applied to the next initiative. Moreover, create internal forums for idea sharing — for instance, an internal blog or demo day where teams showcase AI experiments. A major e-commerce firm's culture



of “open experimentation” led it to host internal innovation jams where employees from any department could propose AI ideas; this surfaced use cases that leadership hadn’t thought of and engaged employees in the transformation (Pupila.ai., The power of) By institutionalising such practices, the company signaled that AI innovation is everyone’s responsibility and opportunity.

Change management

The introduction of agentic AI can trigger anxiety among employees – common fears include job displacement, inability to adapt, or being evaluated by machines. Leadership must tackle these concerns head-on with empathy and clarity across all levels of roles. Engage employees in the AI adoption process – for instance, involve end-users in pilot programs and gather their feedback, so they feel heard and part of the change rather than subjects of it. Providing resources such as an “AI hotline” or an internal chat channel where employees can ask questions about new AI tools can help reduce uncertainty. It’s also effective to share positive stories of employees who have embraced AI and advanced their careers as peer examples can be very persuasive. Other efforts might include workshops on adapting to change, stress management related to AI changes, and clear messaging that while certain tasks will be automated, new opportunities (often more interesting ones) will open.

Lead with ethics

Finally, to foster responsible and confident use of agentic AI, leadership must integrate ethics into the culture. This involves setting and modeling high standards for responsible AI. Make it part of the culture that just because we can automate something doesn’t always mean we should. If an AI agent use case might conflict with the company’s values or customer trust (say, using agents to monitor employees in a way that feels invasive), leaders should be willing to say no. Conversely, leaders should champion agentic AI uses that advance positive outcomes (like fairness in hiring or improving accessibility in products).

Organisation Readiness

Should I build or buy Agentic AI?

As enterprises race to deploy agentic AI, one classic strategy question looms large: Should you build or buy?


For now, we set aside the “partner” option. As the ecosystem is evolving rapidly, most players and their offerings today are still in early stages—too immature for any meaningful co-development. Vendors offer little beyond what can be bought off the shelf and solutions from prominent VC-backed startups aren’t exactly ready to be integrated into existing enterprise environments.

While platform maturity significantly influences the accuracy and reliability of agents across workflows, this not just a technology decision. It requires strategic decision-making that impacts time to value, competitive edge, talent investments, and long-term operating model.


Decision considerations

When formulating an agentic AI strategy, enterprises should break down the decision through several lenses before choosing to build or buy:

Decision lenses	Description
Strategic value of use cases	How core is this AI agent’s capability to your business’s competitive advantage? If the use case involves proprietary processes or IP that could set you apart, an in-house build may offer unique differentiation (linkedin.com, Kartikeya). Conversely, if it’s to deploy in support function (e.g. employee helpdesk), a ready-made solution will suffice.
Time to value	What is your timeline for deployment and ROI? Building from scratch often means a longer development cycle – potentially many months – which delays ROI and competitive benefits. Buying can dramatically accelerate time-to-market as solutions often come pre-trained or ready to deploy, enabling 30–40% faster time to value on average (Gartner)
Cost structure	What does budget and cost horizons look like? In-house development entails significant upfront costs (development effort, infrastructure, training data) and ongoing expenses for maintenance and upgrades. By contrast, buying an solution typically shifts cost to a subscription or usage-based model (OpEx). A caveat here is that integration required for off-the-shelf solutions.
Talent and expertise	Do you have (or can you hire) the AI engineering talent needed? Building advanced AI agents demands skilled data scientists, ML engineers, prompt engineers, etc. – talent that is in short supply and commands high salaries. Retaining such a team long-term is also a challenge. Vendors would often have specialised teams who are experienced in system deployment and servicing. (CIO, Generative AI strategy)




Data privacy and security	How serious are the regulatory, security and data considerations? Building in-house keeps sensitive data internal and allows complete control over data handling. Technical teams often prefer this to avoid sending data to external providers. (Execs in the know) With a buy approach, you must trust the vendor's security controls and policies. Reputable vendors pledge not to use customer data to train their models and many offer data encryption and private cloud options – but due diligence into compliance certifications and ability to support governance requirements is key, especially in regulated industries.
Integration complexity	How easily must the agent integrate with your existing systems and workflows? Building internally gives you freedom to deeply integrate the AI agent into proprietary systems or legacy tools. However, this can be complex – developers may need to connect dozens of APIs, middleware and data sources themselves. By contrast, many vendors provide pre-built integrations and APIs for popular enterprise systems. For example, customer service AI platforms often come with connectors for CRM systems, communication channels (web, SMS, WhatsApp, etc.), and knowledge bases.
Scalability and performance	Anticipate the scale and performance requirements. If your agent needs to handle high volumes (say millions of queries or complex multi-step tasks), can an in-house solution meet that reliably? Building for scale requires robust engineering of infrastructure and model optimisation. Cloud AI vendors offer virtually unlimited scalability on demand – e.g. scaling up GPU instances for heavy workloads – which could be hard to replicate on-premises. (AWS, Cognigy)
Maintenance and evolution	AI agents are not one-and-done software – they require continuous maintenance, tuning, and model updates. Can your teams commit to that ongoing effort? In-house builds mean you carry the full burden of updating models (e.g. adopting the latest GPT model or patching bugs), monitoring performance, and expanding capabilities over time. Vendors, on the other hand, handle updates and improvements as part of their service – e.g. automatically retraining models, pushing new features, content filtering updates, etc. (linkedin.com, Kartikeya).
Use case compatibility	Assess how unique or standard the use case is. This is a crucial lens: If the agent's functionality is fairly standard (like an IT helpdesk assistant that resets passwords or a basic sales email generator), there are likely many vendor solutions or open-source frameworks available. However, if the agent is doing something novel (e.g. a specialised decision-making agent to recommend novel chemical agents for production), existing products might not fit well, tilting toward a custom build.



Navigating vendor landscape and emerging trends

The ecosystem for enterprise AI agents is evolving rapidly. In less than a year, major cloud platforms, enterprise software vendors, open-source communities, and niche emerging startups (Manus, Lovable, Cursor) have flooded the market with new-aged offerings—from pre-trained service bots to orchestration layers that span horizontal applications and job functions. For enterprise leaders, the opportunity is massive—but so is the noise.

- **Cloud Platforms:** The hyperscalers—Microsoft Azure, AWS, Google Cloud, and IBM—provide the foundational infrastructure and models behind many enterprise AI agents. These platforms now offer access to powerful foundation models (e.g. GPT-4, Claude, LLaMA), along with model orchestration tools, governance frameworks, and pre-built agent solutions. Microsoft leads with its Azure OpenAI Service and Power Platform while enhancing its Copilot ecosystem with agentic capabilities and a Copilot Control Panel to enhance governance. (Cloudwars.com) Google Cloud offers Vertex AI and Dialogflow CX for building conversational agents, along with Duet AI and the newly announced Agentspace. (Google Cloud)
- **Enterprise Application Vendors:** CRM and ERP players are increasingly embedding AI agents natively into their ecosystems. Salesforce is at the forefront with Agentforce, a full orchestration layer that lets customers create AI agents inside their CRM workflows. (engineering. Salesforce) ServiceNow augments its Virtual Agent with generative AI for ITSM and employee workflows while Oracle and SAP offer embedded AI agents in HR and ERP suites. These vendors are ideal for enterprises looking to embed agents within their existing application stack.
- **Specialist Agent Vendors and Startups:** The fastest innovation is happening among domain-specific and vertical agent players. Cognigy and Kore.ai offer orchestration tools for conversational AI with pre-built agentic AI skills, 100+ integrations and multi-LLM coordination. (AWS, Cognigy) Moveworks, Aisera, and Yellow.ai specialize in internal support agents (e.g. for IT, HR, legal). Ada, Boost.ai focus on CX and customer service agents that are pre-trained and feature low-code builder modules. Adept AI and AI21 Labs are developing advanced orchestration layers like Maestro, enabling multi-step, tool-using AI agents for business tasks. (ai21.com) These examples highlight only a portion of the evolving landscape.



Emerging vendor trends

The enterprise AI agent space is far from static. Over the next 12–18 months, we expect the following trends to redefine how organisations approach vendor selection:

- **Outcome-Based and Usage-Based Pricing:** Vendors like Salesforce are shifting toward pricing tied to outcomes—for example, charging \$2 per resolved customer service conversation after a usage threshold for its Einstein GPT service. (CIO.com, How will agents be priced)
- **AI Agent Marketplaces:** Salesforce’s AgentExchange now offers 200+ pre-built agent templates and “actions” contributed by partners, while AWS Marketplace hosts deployable agents like Cognigy and Kore.ai. (CRN.com) These marketplaces reduce time-to-value by letting enterprises deploy ready-made solutions for specific domains.
- **Co-Development and Partner Ecosystems:** PwC’s partnership with Microsoft produced Agent OS—a multi-agent orchestration system integrating vendors like OpenAI and Salesforce. Other alliances, like IBM + Cohere or Google + Anthropic, illustrate how integrators and model providers are collaborating to offer scalable, tailored AI agent deployments.
- **Open-Source and Community Accelerators:** Tools like LangChain, Semantic Kernel, and open models like Meta’s LLaMA 2 now allow enterprises to fine-tune agents on proprietary data and build orchestration logic without vendor lock-in. GitHub projects like BabyAGI and Auto-GPT show how quickly new agent capabilities can emerge and be shared.
- **Convergence with Workflow Automation:** AI agents are merging with traditional RPA and BPM tools—UiPath, Zapier and Automation Anywhere are embedding LLMs into bots, while AI vendors now offer orchestration layers akin to workflow engines.
- **Governance, Compliance & Evaluation Standards:** With frameworks like the EU AI Act on the horizon, enterprises are implementing audit trails, content filtering, and bias detection mechanisms. Vendors like Salesforce and Microsoft now embed these controls into their platforms, while buyers increasingly expect vendors to support model transparency, explainability, and enterprise KPIs (e.g. resolution rate, user satisfaction).

Governance and Control

Have I defined the right principles and roles required to govern AI agents?

As Agentic AI systems operate with a high degree of autonomy, robust governance, risk management, and compliance (GRC) frameworks are crucial to ensure trust, security, and ethical behavior. Organisations must evolve traditional AI oversight practices to accommodate the challenges introduced by autonomous agents. This evolution requires a shift from governing individual AI models in isolation to governing AI systems as a whole. As AI deployments become more complex, it's crucial to recognise that AI models don't exist in isolation; they are part of intricate systems that interact with various components, processes, and human operators. Therefore, governance must address these complex interactions to ensure responsible and effective use of Agentic AI.

Foundational principles include accountability (assigning clear responsibility for agent outcomes), transparency and auditability (logging all agent actions and ensuring explainability of decisions), and alignment with business risk (tiering agents by autonomy and impact) (pwc.com, unlocking value with AI agents). To effectively manage these principles, risk assessment must be integrated from the earliest stages of AI system design, embodying a 'compliance by design' mindset. This proactive approach ensures that risk mitigation strategies are developed alongside the AI system, aligning with organisational risk appetite and validating the use case before significant resources are invested. In all cases human oversight is non-negotiable: agents should operate with "a human at the helm for critical decision-making processes", with clear escalation protocols when human judgment is required (pwc.com, unlocking value with AI agents).

Enterprises should establish formal governance bodies to oversee agentic AI. A common model is a central AI Governance Board or Council that sets policies, standards and reviews high-risk projects (gsa.gov, ai compliance plan). This multi-disciplinary body (with representation from IT, legal, risk/compliance, and business units) defines core agentic AI values (fairness, accountability, privacy), formal policies and guidelines, maintains an agent ownership registry (designated owner) and approves new agent deployments. Smaller organisations may combine this with an AI Risk or Safety Team that conducts case-by-case reviews.

Effective agentic AI governance requires a combination of proactive controls and continuous monitoring. A crucial element is the implementation of codified guardrails – modular, reusable components that enforce specific rules and constraints on agent behavior. These guardrails, applied across various use cases, translate high-level policies into actionable rules for data handling, access control, and ethical decision-making. To ensure these guardrails are effective, real-time monitoring of agent activity is essential. This monitoring should encompass performance metrics, quality indicators, and security parameters, with a centralised system providing a holistic view of agent behavior across the enterprise.

Advanced monitoring techniques, such as using LLMs as 'judge models' to evaluate agent outputs and reasoning, can further enhance governance by providing automated oversight of compliance and ethical adherence. Beyond these technical controls, clear roles and responsibilities are vital. AI Risk Officers, AI Product Owners, AI Ethics/Policy Leads, and Model Accountability Leads are crucial for overseeing agent development, deployment, and ongoing management. These roles, along with a RACI matrix, ensure accountability and prevent gaps in governance. Finally, governance workflows must be embedded into the AI agent lifecycle, from policy setting and development to deployment, monitoring, and auditing, to maintain continuous oversight and control.

Governance and Control


What security protocols, tools, and practices are needed for my organisation to build trust towards Agentic AI?

Agentic AI introduces new security challenges that require enterprises to relook at their cybersecurity posture and institutionalise proactive risk identification and rapid response capabilities to guard against a new threat landscape.

AI agents differ from standard applications in that they learn from data and can act autonomously on enterprise systems, amplifying traditional risks. In practice, a compromised agent (e.g. in a supply-chain or logistics function) could issue malicious commands that propagate across the enterprise, causing widespread damage as they serve as gateways to sensitive data and critical systems.

Examples of threat categories that leaders should be aware of:

- **Autonomy Abuse:** AI agents can act independently. If misconfigured or exploited, they may execute harmful actions (e.g., delete data, reconfigure systems) without human approval. Even small bugs or adversarial prompts can escalate into major issues.
- **Unauthorised Access & Privilege Escalation:** Agents often have elevated access to systems and data. If compromised, attackers can impersonate them or trick them into performing high-privilege actions, especially in the absence of strong access controls.
- **Data Exposure & Leakage:** Agents handling sensitive inputs and outputs are vulnerable to prompt injection and hidden payloads. This can lead to unintentional disclosure of confidential data or exfiltration via compromised outputs or memory caches.
- **Model and Data Poisoning (Integrity Attacks):** Adversaries can poison training data, insert malicious documents, or exploit model-loading mechanisms. These attacks degrade performance or create backdoors that activate under specific triggers.
- **Hallucination and Output Manipulation:** Agents may hallucinate content like fake URLs or code packages. Attackers can exploit these hallucinations (e.g., slop squatting) to deliver malware or mislead humans and systems.
- **Prompt Injection and Tool Misuse:** Malicious inputs can override an agent's safeguards via direct or indirect prompt injection. If agents have tool access (e.g., code execution, messaging), attackers can co-opt these tools through deceptive prompts.
- **Multi-Agent Coordination Risks:** In systems with multiple agents, a compromised agent can poison others or hijack shared tasks. Trust assumptions between agents create vulnerabilities that can be exploited for cascading failures.
- **Resource Exhaustion & Denial-of-Service:** Agents can be overloaded with complex tasks or spammed with requests, leading to performance degradation or outages—triggering high cloud costs through excessive compute use.



In response to the new threats, key protocols, tools and governance practices must be embedded across the "Enterprise Data and Tools" layers to ensure safe and responsible deployment:

Core Security Protocols and Tools

These are foundational controls every enterprise should implement to help monitor, constrain, and protect AI agents during operation — especially in high-autonomy or high-risk environments.

- **Strict Access Controls (RBAC/Least Privilege):** Restrict each AI agent's access to the minimum necessary using role- or attribute-based access control. Isolate privileges across zones and enforce strong identity and credentials verification.
- **Sandboxing and Isolation:** Run agents in tightly controlled environments with restricted access to system resources. Consider secure enclaves or containers for high-risk tools like code execution.
- **Immutable Audit Trails and Logging:** Log every agent action and interaction in tamper-evident systems to support real-time monitoring and forensic analysis.
- **Prompt and Input Sanitisation:** Pre-screen all agent inputs to remove hidden instructions or exploit payloads. Only allow validated queries that align with the agent's expected functions.
- **Kill Switches and Human Oversight:** Establish manual override mechanisms for all autonomous agents. High-risk decisions must involve human review until the agent's reliability is well-proven.
- **Agent Firewalls:** Deploy input, data, and trajectory firewalls that enforce prompt hygiene, output filters, and safe behaviour patterns during task execution.
- **Runtime Monitoring Platforms:** Adopt AI runtime security tools to detect threats like data exfiltration or prompt manipulation in real time.
- **Model and Dependency Integrity Checking:** Scan and verify all model weights, libraries, and data sources. Lock down model versions and validate updates in isolated test environments.

Governance and Control

What new expectations will our stakeholders have of us when scaling AI agents?

Stakeholders from customers to employees, regulators, investors, and partners all face new considerations when AI acts autonomously on behalf of a business. As organisations scale their deployment, traditional frameworks of liability, ethics, and oversight become insufficient.

Therefore, organisational accountability must evolve across three critical domains—legal, ethical and operational governance. What is, and will continue to be a technical challenge is now also an accountability issue. Customers, employees, regulators, investors, and partners all expect that businesses deploying autonomous agents remain accountable, transparent, and in control.

Legal responsibility

Organisations bear legal responsibility for actions taken by their AI agents. In most jurisdictions, traditional agency law principles apply: a company remains liable for acts of its AI “agents,” just as it would for human employees (techinsights.linklaters.com).

US regulators like the Federal Trade Commission have explicitly stated that there is “no AI exemption in the books” — companies cannot escape liability for harms or errors caused by AI. Cross-border regulation also looms large: the EU AI Act (and aligned UK regulations) will likely classify agentic AI as “high-risk” or even subject to bans if it manipulates behaviour subliminally (kennedyslaw.com).

Businesses must therefore integrate compliance frameworks that cover both general AI laws (data protection, non-discrimination, safety) and sector-specific rules (e.g. financial regulations, product liability) for any AI-driven function. This spans across regulator mandates, liability exposures and contractual duties. For example, when AI agents are embedded across ecosystems—via vendors, suppliers, or partners—clear contractual terms on data use, model risks, and indemnity become essential. Shared accountability must be codified to avoid ambiguity during disputes.

Ethical imperatives

As agentic AI systems take on greater autonomy, stakeholder expectations around ethical use grow accordingly. Transparency, consent, and fairness become non-negotiable. Customers, employees, and regulators all need to know when an AI agent is acting—and on what basis. (processmaker.com)


Transparency is particularly critical. Agentic decisions must be explainable to those affected, with decision pathways documented and auditable (OECD AI Principles). Users must be informed when they are interacting with an AI system, and what role it plays in shaping decisions that affect them. For instance, in financial services, explainability is now a compliance requirement under both GDPR and proposed EU AI rules. It is therefore key for companies to proactively document how models reach conclusions to build trust among customers while meeting regulatory expectations. (processmaker.com)

Consent and privacy are also under scrutiny. Because agents often operate across sensitive datasets, organizations must ensure proper consent, apply data minimization principles, and implement rigorous access controls. Emerging guidance calls for data governance frameworks that include encryption, anonymization, role-based access, and automated logs (WEF, How Agentic AI will transform).

Fairness must also be continuously assessed. Adaptive agents can unintentionally replicate or amplify social biases from training and testing data. Regular fairness audits, diverse testing datasets, and oversight from ethics boards or AI review committees are becoming best practice. Failure to address bias is not only an ethical failure—it exposes firms to discrimination claims and reputational harm.

Operational governance considerations

Operationally, businesses must treat AI agents like “digital employees” who require supervision, clear objectives, job descriptions and performance metrics. (Protiviti.com). This means specifying each agent’s scope of authority and setting policies (guardrails) around allowed actions with defined boundaries—what data they can access, what actions they can take, and when human intervention is required. All decisions and activities should be logged in tamper-proof audit trails to create traceability and governed under a shared responsibility model between human overseers and agents.



Incident response protocols must be constantly updated to handle agentic AI failures. For example, if an agent makes an incorrect decision affecting a customer—such as denying a loan or issuing a faulty compliance report—the organisation should have a predefined playbook for detection, escalation, correction, and stakeholder communication.


Human oversight remains essential. The principle of “human-above-the-loop” must be preserved in critical functions. Autonomous systems should augment human decision-making, not replace it outright. Enterprise leaders must ensure that all high-stakes decisions remains attached with a human checkpoint.

Internally, agentic AI must be governed like any other enterprise-critical system. Boards and executive teams should assign ownership of AI risk, appoint dedicated roles (e.g. Chief AI Officer), and integrate agent oversight into broader risk management processes. Globally, frameworks like COSO ERM and NIST AI RMF are increasingly being extended to include agent-specific risk scenarios.

Finally, stakeholder assurance must be on a proactive stance. This includes publishing AI accountability reports, disclosing agentic use cases, and engaging investors, customers, and regulators on how risks are being identified, mitigated, and governed. Transparent communication is no longer optional.

As autonomy grows, so must accountability. Organisations can’t rely on legacy structures or trust that ethical intentions will translate into outcomes. Instead, they must build systems that earn trust—through enforceable legal responsibility, embedded ethical design, and rigorous governance practices. Boards must now treat agentic AI as both a business enabler and a risk vector. Accountability isn’t a compliance checkbox—it’s an enterprise-wide competency.

Leading the Way



For leadership teams serious about harnessing Agentic AI, the journey must be intentional, structured, and aligned with enterprise goals. Here's a pragmatic roadmap to move from concept to scaled execution:

1. Assess readiness across the enterprise

- Audit your data, talent, process maturity, and governance baseline.
- Identify critical gaps—especially in infrastructure (e.g., agent-ready data platforms) and change readiness.

2. Align leadership around strategic outcomes

- Brief the C-suite and business unit leaders on the transformative potential of Agentic AI.
- Anchor early initiatives to specific business priorities—cost, speed, revenue, customer experience.

3. Establish governance and risk frameworks

- Establish an AI Governance Board and Ethics Committee.
- Define roles, ownership models, accountability protocols, and guardrails for agent autonomy and risk., accountability protocols, and security standards.

4. Run focused pilot programs

- Use “Agent Designer” tools to prototype in specific functions (e.g., sales, customer service, compliance).
- Prioritise use cases that are structured, measurable, and repeatable.

5. Build cross-functional teams

- Form agile squads combining business leads, data engineers, AI developers, and risk managers.
- Foster collaboration using common toolkits, operating environments and shared Agent Platforms.

6. Scale what works

- Use pilot results to inform enterprise-wide rollout.
- Expand agent deployment across operations, customer service, R&D, and supply chains.

7. Monitor, Audit, and Evolve

- Set up continuous evaluation cycles with both short- and long-term performance metrics.
- Regularly audit agent behaviour for compliance, ethical alignment, and optimisation opportunities.

By following this roadmap, organisations can transition from experimentation to enterprise-wide transformation, unlocking the full competitive advantage that Agentic AI offers

Conclusion



Agentic AI marks the beginning of a structural and cultural shift in how enterprises operate, compete, and lead. As this whitepaper has outlined, it is not just another technology to integrate—it is a test of executive clarity across four foundational pillars: strategic direction, business value, organisational readiness, and governance.

The central insight is clear: agentic AI changes the nature of execution due to increased autonomy, improved learning and collaboration between agents and humans. Emerging use cases will no longer build on top of systems that support decisions but increasingly leverage systems that make them. That change demands more than pilots and prototypes—it demands a reinvention of how leaders define value, design organisations, and distribute accountability. Agentic systems don't just optimise what we already do. They promise to enable what wasn't previously possible.


To realize the potential of this wave of innovation with AI agents, leaders must set direction before agents set the pace. They must build teams that work with this new type of AI, not around it. They must prove impact early—through productivity, speed, and resilience—that is well clarified and documented at the domain level. And they must maintain critical oversight on processes as autonomy scales, ensuring trust is built into every decision loop even when the roles of human change in significant ways.

This whitepaper has offered a structured foundation to begin that journey—twelve questions that brings clarity, challenge assumptions, and surface the trade-offs that matter. Still, this is only the starting point. Further research to explore how agentic AI will move deeper into the enterprise, its organisational and societal implications is needed—not just on technical capabilities, but on labor markets, governance models, trust frameworks, and long-term ethical and societal impact.

What's clear is this: leaders who treat agentic AI as a narrow IT project will fall behind. Those who see it as a company-wide transformation—of structure, skills, culture, and responsibility—will define the next frontier. The playbook is emerging. And the leadership test has already begun.

Before you commit to the next deployment, bring your leadership team together and ask: are we aligned on what matters most? You've read the questions and our suggestions—now make them yours.

Endnotes



<https://partners.wsj.com/dell-and-nvidia/the-time-is-how-for-ai/the-rise-of-agentic-ai/>

https://www.mulesoft.com/lp/reports/connectivity-benchmark?_gl=1*1au445d*_gcl_au*NzQ2OTIyODY2LjE3NDIzNTEINDM.*_ga*Nzk0ODkxMjg5LjE3NDIzNTEINDM.*_ga_3VHBZ2DJWP*MTc0NDklnjY4MS43LjAuMTc0NDklnjY4MS42MC4wLjA&_ga=2.122052683.178839360.1744944112-794891289.1742351543

https://www.adamsstreetpartners.com/wp-content/uploads/2025/03/2025-03-12-The-Rise-of-Agentic-AI_P.pdf

<https://hbr.org/2024/12/what-is-agentic-ai-and-how-will-it-change-work>

<https://developer.nvidia.com/blog/building-ai-agents-to-automate-software-test-case-creation/>

<https://arxiv.org/pdf/2504.15474>

<https://www.supplychainbrain.com/blogs/i-think-tank/post/40728-price-smart-act-fast-agentic-ais-role-in-retails-next-chapter>

<https://www.ark-invest.com/articles/analyst-research/ai-agents-could-transform-enterprise-spending#:~:text=Early%20examples%20of%20the%20benefits,and%20marketing%20expenses%20and%20customer>

https://cdotimes.com/2025/01/06/ups-and-agentic-ai-a-case-study-in-logistics-innovation/#:~:text=_edge%20in%20the%20logistics%20sector

https://www.moveworks.com/content/dam/moveworksprogram/pdfs/customer-stories/Moveworks_CaseStudy_Power-Design.pdf#:~:text=Results%20-%20Automated%201%2C000%2B%20hours,Desk%20of%20the%20Year%20by
<https://www.weforum.org/stories/2016/07/google-harnesses-the-power-of-ai-to-cut-energy-use/#:~:text=Google%20said%20it%20has%20been,of%20a%20DeepMind%20AI%20system>

<https://www.walturn.com/insights/the-cost-of-implementing-ai-in-a-business-a-comprehensive-analysis#:~:text=maintenance,term%20AI%20success>

https://www.linkedin.com/posts/appenz_a-number-of-agentic-ai-services-are-now-being-activity-7292307575973285888-gcRN/

<https://www.siroccogroup.com/demystifying-agentic-ai-pricing-what-to-consider-when-evaluating-different-pricing-models/>

<https://ai-infrastructure.org/the-hidden-costs-challenges-and-tco-for-gen-ai-adoption-in-the-enterprise-sept-2023/#:~:text=But%20we%20also%20found%20that,subscription%20pricing%20for%20cloud%20models>

<https://www.cio.com/article/3624540/how-will-ai-agents-be-priced-cios-need-to-pay-attention.html>

<https://www.linkedin.com/pulse/crunching-numbers-cost-analysis-ai-agents-enterprise-ashoori-phd-kp7ve/>

<https://www.ibm.com/think/insights/ai-economics-compute-cost>

<https://www.mckinsey.com/industries/financial-services/our-insights/insurance-blog/the-potential-of-gen-ai-in-insurance-six-traits-of-frontrunners>

<https://www.salesforce.com/blog/vector-database/#:~:text=This%20is%20critical%20for%20the,social%20media%20posts%2C%20making%20it>


<https://www.cio.com/article/3808569/knowledge-graphs-the-missing-link-in-enterprise-ai.html#:~:text=and%20subject%20matter%20experts%20who,what%20the%20data%20actually%20means>

<https://medium.com/@jerrymartejr/feature-stores-the-key-to-bridging-ai-and-data-engineering-3ae24834c1l4#:~:text=integration%20of%20features%20across%20training>

<https://enterprisesight.com/the-agentic-ai-era-fueled-by-knowledge-graphs/#:~:text=Agentic%20AI%20systems%20depend%20on,context%20at%20the%20right%20time>

<https://www.workato.com/the-connector/agentic-orchestration-future/>

Endnotes



<https://www.cio.com/article/3541227/generative-ai-strategy-dilemma-buy-build-or-partner.html#:~:text=Yet%20GenAI%20technologies%20and%20processes, on%20your%20chosen%20use%20cases>

<https://execsintheknow.com/blog/building-vs-buying-ai-agents-for-your-contact-center-why-you-no-longer-have-to-choose/#:~:text=Nothing%20is%20more%20important%20to, leading%20to%20inadvertent%20data%20leaks>

<https://aws.amazon.com/marketplace/pp/prodview-6c25zypcl2fro#:~:text=Highlights>

<https://cloudwars.com/ai/pwc-ai-agent-provides-orchestration-oversight-of-multi-vendor-agent-deployments/#:~:text=While%20Microsoft%20has%20its%20tightly, SAP%2C%20Salesforce%2C%20Workday%2C%20and%20more>

<https://engineering.salesforce.com/engineering-agentforce-building-a-modular-multi-model-framework-for-enterprise-ai-agents/#:~:text=The%20mission%20with%20Agentforce%20is, for%20the%20swapping%20of%20models>

<https://www.ai21.com/knowledge/ai-agent-frameworks/#:~:text=Introducing%20AI21%20Maestro%3A%20An%20AI, solve%20complex%20tasks%20Learn%20More>

<https://www.crn.com/news/ai/2025/salesforce-launches-ai-agents-marketplace#:~:text=The%20marketplace%20launches%20with%20more, through%20Thursday%20in%20San%20Francisco>

<https://cloud.google.com/products/agentspace?hl=en>

<https://www.cio.com/article/3972714/ms-ai-agents-become-teammates-a-new-enterprise-form-emerges.html>

<https://time.com/charter/7281015/ai-agents-and-what-comes-after-the-org-chart/>

<https://www.pupila.ai/news/the-power-of-collaboration-in-the-ai-era-unlocking-innovation-through-partnerships#:~:text=operating%20in%20silos%20but%20are, together%20to%20drive%20AI%20innovation>

<https://www.linkedin.com/pulse/ai-agents-changing-enterprise-org-charts-are-you-ready-fpy4e/>

<https://rossdawson.com/framework-agentic-ai-core-patterns-for-organization-design/#:~:text=In%20essence%2C%20successful%20organizational%20redesign, organizations%20seeking%20to%20navigate%20this>

<https://www.deloitte.com/lu/en/our-thinking/future-of-advice/agent-ai-the-new-frontier-in-ai-evolution.html>

<https://www.ciodive.com/news/AI-talent-upskilling-accenture/742741/#:~:text=%2A%20Nearly%20two, the%20skills%20their%20organization%20needs>

<https://www.cio.com/article/3966870/how-it-leaders-use-agentic-ai-for-business-workflows.html#:~:text=billion%20or%20more%20in%20revenue, top%20perceived%20benefit%20of%20it>

<https://www.weforum.org/stories/2025/01/the-trust-imperative-5-levers-for-scaling-ai-responsibly/#:~:text=, evolving%20regulatory%20landscape>

<https://www.pupila.ai/news/the-power-of-collaboration-in-the-ai-era-unlocking-innovation-through-partnerships#:~:text=Open%20experimentation%20and%20idea>

<https://www.ciodive.com/news/AI-talent-upskilling-accenture/742741/#:~:text=Accenture%20deployed%20reskilling%20as%20part, employees%20collaborate%20with%20AI%20agents>

<https://hbr.org/2025/03/employees-wont-trust-ai-if-they-dont-trust-their-leaders>

<https://www.pwc.com/us/en/tech-effect/ai-analytics/responsible-ai-agents.html#:~:text= detect%20unauthorized%20access%20or%20anomalies>

<https://www.gsa.gov/technology/government-it-initiatives/artificial-intelligence/ai-guidance-and-resources/ai-compliance-plan#:~:text=These%20governance%20bodies%20promote%20the, for%20all%20AI%20use%20cases>

<https://www.modelop.com/blog/whos-accountable-in-the-enterprise-for-ai-and-its-risks#:~:text=Given%20the%20pivotal%20role%20of, for%20issues%20that%20absolutely%20will>

<https://yu-ishikawa.medium.com/the-rise-of-ai-agents-in-the-enterprise-part-2-designing-an-enterprise-agent-governance-framework-e3b1f0ba950f#:~:text=I, analogous%20to%20a>

Endnotes



<https://thoropass.com/blog/compliance/ai-grc-team/#:~:text=Data%20Protection%20Officer%20%2F%20Chief,Information%20Security%20Officer>

<https://www.credo.ai/recourseslongform/from-assistant-to-agent-navigating-the-governance-challenges-of-increasingly-autonomous-ai#:~:text=Oversight%20and%20Governance%20Structures%3A%20Given,reviewed%20by%20a%20human%20promptly>

<https://techinsights.linklaters.com/post/102k5ac/understanding-agentic-ai-the-power-of-autonomy#:~:text=In%20terms%20of%20risk%2C%20as,AI%20exemption%20in%20the%20books>

<https://kennedyslaw.com/en/thought-leadership/article/2025/agentic-ai-what-businesses-need-to-know-to-comply-in-the-uk-and-eu/#:~:text=Businesses%20must%3A>


<https://www.processmaker.com/blog/ethical-considerations-of-agentic-ai/#:~:text=Transparency%20and%20Explainability>

<https://www.weforum.org/stories/2024/12/agentic-ai-financial-services-autonomy-efficiency-and-inclusion/#:~:text=,of%20defining%20responsibility%20and%20liability>

https://www.protiviti.com/sites/default/files/2025-04/protiviti_newsletter-bp186-agentic-ai_global.pdf#:~:text=To%20mitigate%20these%20risks%20and, and%20guardrails%20for%20the%20agent's

https://www.galileo.ai/blog/ai-agent-metrics#:~:text=_delivery%20or%20improved%20customer%20experience

<https://www.iaitransformation.com/insights/what-are-the-key-metrics-for-measuring-ai-governance>



Contact Us

Laura Flores
laura@wahspark.com

Spark
www.wahspark.com