




The business imperative for Agentic AI

March 2026

Table of contents


Introduction	3
• Looking beyond Generative AI benefits	4
• Luxembourg's next AI leap: Agentic AI	5
• Agentic AI powering transformations across various business functions	6
Key considerations for undertaking a successful Agentic AI journey	8
1. Which business processes best suit Agentic AI implementation, and what criteria should be used to evaluate their sustainability?	9
2. Is your technology ecosystem ready for Agentic AI? What foundational components are essential for enabling agentic workflows?	11
3. How do you define and measure the success of Agentic AI interventions, and what expectations should be set within the organisation?	13
4. What is the most effective approach to building, scaling and sustaining an Agentic AI journey that aligns with your organisation's unique goals and capabilities?	15
5. What's needed to equip and empower employees to thrive alongside AI agents, as they become an integral part of the workforce?	18
6. What are the key considerations for ensuring the responsible deployment of Agentic AI?	21
Conclusion	23



Introduction

As the demand for automation grows to transform business processes with complete autonomy, organisations are shifting beyond Robotic Process Automation (RPA) and embracing advanced technologies such as Agentic AI. While both aim to improve efficiency, they operate on fundamentally different principles. RPA automates discrete, rule-based tasks and requires human intervention for exceptions, logging or decision-making. It is effective for repetitive, structured processes but lacks adaptability and autonomy.

Agentic AI, on the other hand, represents a significant evolution. It is goal-oriented, context-aware and capable of autonomous decision-making. These intelligent agents execute tasks and validate, correct and optimise operations in real time. Suppose a task cannot be completed at the first pass. In that case, the agent escalates it to a supervisory agent, only involving humans where explicitly required, such as for compliance checks or strategic approvals. This makes Agentic AI far more resilient, adaptive and efficient than traditional automation.



AI agents can orchestrate complementing tasks across these systems

Use case

- RPA: Collects supplier delivery history, product quality and cost data
- ▼
- GenAI: Identifies underperforming suppliers and selects more cost-effective alternatives
- ▼
- Agentic AI: Initiates contract renewal or negotiates agreements with new suppliers

Source: Deloitte Insights


Looking beyond Generative AI benefits

Before the beginning of 2025, the landscape was dominated by discussions surrounding AI and Generative AI (GenAI). Looking back, 2023 stands out as a pivotal year for GenAI, capturing worldwide attention and showcasing AI's remarkable creative and cognitive capabilities in various fields, including text generation, imagery creation and coding. By 2024, the narrative shifted from exploration to execution, with enterprises embedding GenAI into core business processes and workflows¹. However, GenAI's limitations, particularly in reasoning, decision-making and autonomous execution, have paved the way for the emergence of Agentic AI systems.

¹<https://www2.deloitte.com/us/en/pages/consulting/articles/state-of-generative-ai-in-enterprise.html>

GenAI

Agentic AI



Source: Deloitte Insights

Luxembourg's next AI leap: Agentic AI

In recent years, the focus in Luxembourg has begun to shift from organisations integrating Generative AI into core business processes toward the emergence of **Agentic AI systems**. Within the country's highly digitalised, regulated and data-intensive ecosystem, enterprises are increasingly exploring autonomous AI agents to manage growing operational complexity while maintaining trust and regulatory compliance.

Luxembourg's mature digital foundations — including robust financial market infrastructures, advanced data platforms, and established digital identity and trust frameworks — have contributed to a steady increase in data volumes and process interdependencies. This evolving environment provides a strong

basis for the adoption of Agentic AI across key sectors such as financial services, insurance, asset management and public administration.

Agentic AI enables organisations to improve operational efficiency by automating front-, middle-, and back-office functions, including customer engagement, risk and compliance activities, finance operations, and reporting. Crucially, this development goes beyond traditional task automation. Agentic systems are designed to operate autonomously and adaptively within clearly defined governance and regulatory guardrails, supporting more informed and responsive decision-making across organisational functions.

Agentic AI powering transformations across various business functions

Agentic AI enables businesses to operate efficiently by automating the front, middle and back-office functions, including customer outreach and engagement, product innovation and finance operations. What's unique and interesting here is that the shift is not just about task automation; it is about empowering systems to act autonomously and adaptively, driving smarter outcomes

across functions. For example, Agentic AI will empower retail systems to act autonomously and adaptively, optimising inventory in real time, dynamically adjusting pricing based on demand and competition, personalising customer engagement across channels and streamlining supply chain operations. This will ultimately drive smarter outcomes across the retail value chain.

Agentic AI's applications across the retail and FMCG value chain

Research and development	Manufacturing	Sales and marketing	Supply chain management	Retailer and distribution	Enabling areas
Autonomous product innovation	Autonomous quality control	Multi-channel campaign orchestrator	Demand forecasting agent	Shelf stocking assistant	Autonomous finance analyst
Smart experimentation	Predictive maintenance agent	Dynamic pricing agent	Inventory balancer	Planogram compliance checker	HR talent scout
Patent landscape navigator	Production scheduling optimiser	Sentiment-driven brand manager	Logistics route optimiser	Retail forecast collaboration agent	Compliance and risk monitoring agent
Consumer sentiment synthesiser	Energy efficiency advisor	Autonomous content creator	Supplier risk monitor	Autonomous retail order management	Cybersecurity sentinel
Sustainability optimiser	Waste reduction agent	Retailer-specific offer optimiser	Autonomous procurement agent	Autonomous sales beat planning	ESG reporting agent
AI-driven competitive analysis	AI-driven process optimisation	Autonomous customer segmentation	Disruption detection and mitigation agent	Geo-targeted promotion agent	Legal contract review agent

Source: Deloitte Insights

Agentic AI's applications across the FSI sector value chain

Research and development	Finance	Sales and marketing	Customer service	Operations
Product innovation agent	Under-writing recommendation agent	Cross-sell recommendation agent	Autonomous case resolution agent	Loan processing agent
Risk modelling agent	Risk reasoning assistant	Autonomous promotions campaigns	KYC/Onboarding compliance agent	Claims adjudication agent
Regulatory intelligence agent	Dispute resolution agent	AI sales advisor agent	Voice of customer agent	Fraud detection agent
Behavioural insights agent	Expense anomaly agent	Dynamic pricing agent	Grievance escalation agent	Payment reconciliation agent
Investment strategy generator	Regulatory reporting agent	Competitor intelligence agent	Proactive service agent	Branch/Channel performance agent

Source: Deloitte Insights



Key considerations for undertaking a successful Agentic AI journey

As enterprises increasingly embrace the transformative potential of Agentic AI, a strategic and well-structured approach becomes essential to adopt and realise its full outcome value.

Six critical questions every enterprise must answer to harness the full power of Agentic AI

1. Which business processes best suit Agentic AI implementation, and what criteria should be used to evaluate their sustainability?
2. Is your technology ecosystem ready for Agentic AI? What foundational components are essential for enabling agentic workflows?
3. How do you define and measure the success of Agentic AI interventions, and what expectations should be set within the organisation?
4. What is the most effective approach to building, scaling and sustaining an Agentic AI journey that aligns with your organisation's unique goals and capabilities?
5. What's needed to equip and empower employees to thrive alongside AI agents, as they become an integral part of the workforce?
6. What are the key considerations for ensuring the responsible deployment of Agentic AI?



1. Which business processes best suit Agentic AI implementation, and what criteria should be used to evaluate their sustainability?

Agentic AI transforms automation from a static tool into a collaborative digital workforce. As it is reasoning-based and proactive, organisations must carefully select complex, goal-driven teams that can benefit from autonomy and continuous learning to unleash their full potential. Implementing agentic AI requires a thoughtful approach; it should not be a generic

solution adopted for its own sake. Instead, organisations must carefully identify high-impact use cases that align with their strategic objectives. By doing so, they ensure that Agentic AI is more than a technological upgrade; it becomes a strategic enabler of business transformation.

Therefore, it is critical to identify the right processes and use cases that can be agenticised. Agentic AI is not a one-size-fits-all solution; only select processes with the right characteristics can truly benefit from its capabilities. Organisations must focus on use cases that deliver meaningful influence quickly and sustainably rather than pursuing implementation for novelty or experimentation alone.

To assess suitability for Agentic AI, organisations should rigorously evaluate processes against key criteria: high complexity and context dependence, a clear need for autonomous decision-making, goal-oriented workflows, multi-step and interdependent tasks and potential for continuous learning and optimisation. These attributes help distinguish processes that can be effectively agenticised and drive real business value. The table below outlines these criteria in detail.

Criteria to evaluate if a use case is Agentic

Criteria	Description	Ideal vs. non-ideal examples of Agentic processes	
Reasoning, complexity and context dependence	Processes that require logical reasoning, decision-making and adaptation to changing inputs are ideal	Ideal - Customer service handling, parsing and sentiment - Supply chain optimisation - Financial forecasting planning	Non-ideal - Forecast model building - Customer segmentation
Need for autonomy and escalation	Tasks that involve following instructions require underlying agents to operate independently, demonstrate autonomy in reasoning and escalate when needed	Ideal - IT incident management - Compliance monitoring - SDLC process automation	Non-ideal Code generation as a step only
Process with a logical end	Workflows that have clearly defined end goals and do not end at an intermediate step. A use case might be a part of a bigger process and should independently have an end goal	Ideal Expense verification	Non-ideal Digital Twins of finance processes
Action and goal-oriented workflows	Tasks that focus on achieving outcomes rather than simply execution. The use case under consideration may involve reasoning and planning actions, but may not necessarily perform specific business actions	Ideal - Resolving a customer issue - Completing a procurement cycle	Non-ideal - Sending trigger-based pre-drafted communications - CRM entry creation from call centre leads
Multistep and interconnected tasks	Workflows involving multiple steps, tools, or systems. All steps to execute a process may (or may not) involve AI; sub-agents can include automation agents, GenAI based respondent agents, deep learning-based matching agents, etc.	Ideal - Customer onboarding - Claims processing - Project coordination	Non-ideal Document comparison as a point solution
Cyclic and repetitive process	Workflows that have a definitive cycle period and repeat frequently	Ideal CV qualification	Non-ideal Attrition Causality and actions
Non-explanatory workflows	A problem that does not require causality and explanations. For instance, Agentic AI cannot answer questions like "Why is revenue stagnant?"	Ideal Change request management	Non-ideal Summarisation on a dashboard
Continuous learning and optimisation potential	Processes that improve over time with feedback	Ideal - Marketing campaign management - Fraud detection	Non-ideal Email list segmentation based on fixed rules



2. Is your technology ecosystem ready for Agentic AI? What foundational components are essential for enabling agentic workflows?

Depending on the maturity of their technology ecosystem, organisations can be at varying stages of readiness for Agentic AI adoption. For instance, organisations already using GenAI, automation and Large Language Models (LLMs) typically operate in a multisystem environment with robust data pipelines, orchestration layers and integration frameworks. These organisations often have a preferred LLM embedded into their workflows, supported by automation platforms (RPA and iPaaS) and APIs that enable seamless interaction across systems. Their environments are generally open and modular, allowing for rapid experimentation and deployment of agentic workflows rather seamlessly.


In contrast, organisations not yet using GenAI or automation face a steeper path to Agentic AI implementation. Their systems may be siloed, lacking interoperability and real-time data exchange, limiting autonomous agents' effectiveness. Without foundational model or a chosen LLM, these environments require significant groundwork, such as modernising infrastructure, integrating systems and establishing governance frameworks, before Agentic AI can be meaningfully deployed. Agentic AI adoption is more incremental for these organisations,

often starting with pilot use cases in isolated domains before scaling across the enterprise.

For organisations at any stage of readiness, it is critical to take a long-term view when defining their technology roadmaps, especially as Agentic AI begins to reshape enterprise capabilities. Agentic AI systems are more advanced versions of traditional automation or GenAI and represent a paradigm shift towards autonomous, goal-driven agents operating across dynamic environments with minimal human intervention.

This shift demands a rethink of the overall technology strategy. Organisations must assess their current infrastructure, identify automation, data flow and governance gaps, and invest in platforms supporting continuous learning, agent coordination and ethical oversight. Those proactively aligning their tech ecosystems with the requirements of Agentic AI will be better positioned to scale these capabilities enterprise-wide and harness transformative value over time.

Agentic AI readiness checklist



If the organisation is already using automation or GenAI

Requirements to be Agentic AI ready (Non-exhaustive)

- Ensuring hyperscaler readiness
- Ability to host serverless agent workloads (Lambda, Kubernetes, etc.)
- Seamless data access layer with APIs, data lakes
- Ability to expose data to APIs, which would be linked to agents for back-and-forth data transfer
- Secure API gateway
- Framework support for multi-agentic workflows
- Availability of vector databases or RAG with session-based long-term memory
- Integration with logging/monitoring tools
- CI/CD pipelines for agents



If starting from scratch

Requirements to be Agentic AI ready (Non-exhaustive)

- Building in-house tech stack (Cloud, APIs, Vector DBs, Orchestrations tools, etc.)
- Investment in middleware or iPaaS to allow agents to interact with systems
- Building RESTful APIs for core systems
- Adopting an API gateway
- Establishing a data lake or warehouse
- Setting up CI/CD for agents
- Deploying LLMs for tools
- Enforce audit logs, sandboxing



3. How do you define and measure the success of Agentic AI interventions, and what expectations should be set within the organisation?

The success of Agentic AI implementation hinges on how well it delivers against the organisation’s ultimate business outcomes and strategic goals. These goals vary widely from automating manual tasks to reducing costs and improving efficiency to enabling intelligent agents that adapt to dynamic environments and drive innovation.

Before embarking on an Agentic AI journey, organisations must clearly define what they aim to achieve. This clarity ensures that AI efforts are aligned with broader business objectives, enabling meaningful measurement of success and return on investment. Organisations risk deploying AI without purpose and well-defined goals, leading to wasted resources, unmet expectations and minimal business impact. Clear goals also guide the data strategy, ensuring the correct data is collected and used to effectively train and optimise AI models.

Importantly, Agentic AI is not a point solution like many GenAI tools. It is process-oriented, meaning it must be deeply integrated into business workflows and continuously evolve with them. This requires a significant upfront investment in planning,

cross-functional collaboration and change management. Agentic AI systems must be trained, tested and refined over time, with feedback loops and performance monitoring built into the implementation.

Organisations must approach Agentic AI as a long-term strategic initiative rather than a quick fix. Its success depends on the technology and the organisation’s readiness to adapt processes, align teams and commit to continuous improvement.

Measuring the success of Agentic AI requires a more nuanced approach than traditional automation. While RPA is often evaluated using binary metrics, such as time saved, errors reduced or tasks completed, Agentic AI introduces cognitive complexity, reasoning, multi-step decision-making, continuous learning and context retention. Success must be assessed across dimensions, including reasoning accuracy, decision autonomy, exception handling, customer experience improvements, efficiency gains, multilingual and multimodal capabilities, and overall convenience. These metrics reflect the broader and more dynamic role Agentic AI plays in business transformation.

Agentic AI value drivers

Business goal	Value-driver	How can Agentic AI drive these results
Enhancements to the “as-is process”	Faster Turn Around Time (TAT)	Streamlines operations with intelligent task routing and automation
	Reduced error rates	Continuously learns from data to improve accuracy and reduce rework
	Improved efficiency and process accuracy	Identifies inefficiencies and automates low-value tasks
	Optimised future-state process flows	Simulates and recommends efficient workflows using predictive modelling
	Improved profitability	Enhances operational margins through automation and intelligent insights
	Enhanced utilisation	Identifies underused assets and reallocates them effectively
Risk management	Comprehensive risk assessment and mitigation	Detects anomalies and predicts risks using advanced analytics
	Feedback loop for continuous improvement	Integrates feedback mechanisms to refine risk models and responses
Business impact	Revenue generation	Drives top-line growth through better customer engagement and operational excellence
	Cost avoidance	Proactively identifying and mitigating potential future expenses
	ROI	Measures and maximises ROI through targeted AI initiatives aligned with strategic goals



4. What is the most effective approach to building, scaling and sustaining an Agentic AI journey that aligns with your organisation's unique goals and capabilities?

Agentic AI pilots have shown promise; however, organisations need to think about how they can build and scale agentic AI solutions to deliver transformational business value. The first step for any organisation envisioning the implementation of Agentic AI is to carefully evaluate the strategic approach,

considering whether to build, collaborate or adopt a hybrid approach for implementing Agentic AI solutions. Additionally, they must assess and decide when to use RPA, LLM, single agent or multi-agent capabilities.

Three key strategic approaches to implement Agentic AI solutions:

- **Build:** Developing Agentic AI solutions in-house
- **Partner:** Collaborate with Agentic AI specialist partners to build end-to-end solutions
- **Hybrid (Buy + Partner Support):** Purchase pre-built generic AI agents or platforms + Use partner for implementation

Build vs Partner vs Hybrid approach

Criteria	Build	Partner	Hybrid
Strategic value	When Agentic AI is core to your IP, product, or long-term strategy	When Agentic AI is important but requires high customisation and domain-specific capabilities	When Agentic AI is not core, off-the-shelf solutions will suffice the requirement
Data protection	Data security is a significant ethical concern	Data can be shared with external partners	Only limited data can be shared
Talent availability	Strong in-house data/AI teams available	Some expertise but seek external augmentation	Limited AI/tech expertise internally
Time-to-market pressure	High upfront investment; better long-term ROI at scale	Shared cost/risk model with scalable roadmap	Instant deployment and quick time-to-value are top priorities
Cost and scalability	Longer lead time, full control	Collaborative model that allows for co-investment, minimising risk while maximising value	Prefer solutions with set budgets and implementation schedules

Risks and rewards

- **Building on own:** Developing in-house Agentic AI solutions may promise complete control, but it is often the costliest and the riskiest path. The intense global competition for AI talent and the breakneck pace of technological advancement make it difficult for most organisations to keep up without significant, sustained investment.
- **Partner:** Collaborations are necessary for organisations building long-term value from Agentic AI. However, effective strategic Agentic AI alliances work differently from traditional vendor relationships. Agentic AI technology is still rapidly evolving, and building and implementing it is complex, requiring closer collaboration and higher degrees of trust.
- **Hybrid approach:** This approach enables faster time-to-value, especially for standard use cases such as customer support automation, document processing or predictive analytics. It also allows organisations to build long-term internal capabilities, which will prove to be the strategic value driver in the long term. However, the trade-off lies in limited customisation, vendor lock-in and reduced control over the AI roadmap.

Furthermore, organisations must carefully evaluate their processes, choosing between RPA, LLMs, single agents and multi-agent systems that depend on the task’s complexity,

adaptability and decision-making requirements. The more effective the choice, the more ROI it generates.

When to use an RPA vs LLM vs Single Agent vs Multi Agent?

	No reasoning	Self-reflection/ Correction	Plan and execute (ReACT)	Collaboration/ Teams
	Systems operating without decision-making, performing pre-programmed tasks	Systems perceiving, evaluating and adjusting their actions for continuous improvement	Advanced systems planning and executing tasks effectively	Systems uniting and synchronising tasks for achieving a shared goal
Semi-autonomous tasks				
Tasks requiring human aid, or additional resources and not completely “automatable”	NA	LLM w/memory	LLM w/memory	Multiple LLMs w/ memory
Single executable tasks				
Task with a single goal and output	RPA bot	Single agent	Single agent	Multi-agent
Workflow of tasks				
A series of interconnected tasks with individual actions performed in a sequence or in parallel to achieve a larger goal	RPA workflow	Multi-agent w/supervisor	Multi-agent w/supervisor	Multi-agent w/ supervisor

Source: Deloitte Insights





5. What's needed to equip and empower employees to thrive alongside AI agents, as they become an integral part of the workforce?

Agentic AI is changing how people work and the nature of their jobs, creating profound implications for employment. The future of work will shift beyond automation. It introduces intelligent, autonomous systems capable of decision-making, learning and collaboration that are critical. To harness the full potential of Agentic AI, businesses must prioritise:

1. Upskilling employees with requisite skills for Agentic AI
2. Redefining job roles to emphasise human creativity and strategic thinking
3. Preparing leadership to orchestrate the future of human-Agentic AI collaboration



1. Upskilling employees with requisite skills for Agentic AI

Intelligent agents will increasingly handle routine operations, enabling the workforce landscape to evolve. To work with agentic AI systems, employees must possess skills that help systems improve their performance.

Key requisite skills for the Agentic AI future

- **Data scientists:** Agentic AI thrives on intelligence, but it is data scientists who bring it to life. From designing solutions to training models and driving continuous performance, they are the architects behind every successful AI agent. Without them, Agentic AI is just potential, not power. The core skill includes transforming raw data into actionable insights to support agentic reasoning, dynamic decision-making and personalised interactions.
- **Machine Learning (ML) engineers:** The rise of Agentic AI, where autonomous agents can reason, learn and act with minimal human input, is fuelling a sharp increase in the demand for ML engineers. The core skills include building scalable learning pipelines, models and infrastructure that power autonomous agent behaviours across varied environments.
- **Prompt engineers:** As enterprises adopt Agentic AI systems composed of intelligent agents that reason, plan and act autonomously, the role of prompt engineers has become crucial. Their core skills include designing prompt structures and scaffolding mechanisms to guide LLM agents towards reliable, safe and context-aware behaviour.
- **AI architects:** They lead the end-to-end design of Agentic AI systems, integrating reasoning loops, planning capabilities, tool use (via APIs or environments) and multi-agent coordination. Their core skills include designing and governing the architecture of agentic systems, ensuring coherence, safety, interoperability and human alignment.
- **LLM engineers:** Engineering systems for persistent or episodic memory in agent conversations is a crucial part of the Agentic AI model. It requires specialisation in customising and optimising foundation models to function as reliable, controllable agents across tasks such as reasoning, planning and tool use.
- **AI Ops engineers:** Operationalisation, monitoring and lifecycle management of autonomous agents on a regular interval is essential in real-world environments. Core compatibility includes tracking agent behaviour, prompt drifts, hallucinations and usage metrics. This includes automating the deployment and updates of autonomous agents across environments. They also manage long-term memory stores, embeddings and semantic retrieval for agents.



2. Redefining job roles to emphasise human creativity and strategic thinking

Roles include expectations for working alongside agents, validating AI decisions and steering outcomes through critical thinking.

- **Rise of hybrid roles:** With the growth of Agentic AI, new hybrid roles, such as AI collaboration specialist, conversational AI specialist and agent supervisor, are expected to be created, blending technical understanding with domain expertise and judgment.
- **Focus on creativity and strategy:** Employees need to spend more time on tasks involving creativity, problem framing and innovation, enabled by intelligent agents automating routine workflows.
- **From execution to orchestration:** Employee roles are shifting from executing discrete tasks to orchestrating end-to-end outcomes, using agents as co-pilots or assistants to drive productivity, personalisation and business impact



3. Preparing leadership to orchestrate the future of human-Agentic AI collaboration

As intelligent agents become embedded in day-to-day workflows, leaders must be trained to manage hybrid teams, where humans and intelligent agents collaborate to drive outcomes. This shift demands a new kind of leadership that blends digital fluency with empathy, adaptability and responsible AI oversight.

- **Culture and change management:** Leaders should create environments that empower employees to collaborate confidently with intelligent agents and add meaningful human insight. They should embrace experimentation, innovation and learning through trial and error.
- **Prioritise human outcomes beyond metrics and efficiency:** As Agentic AI takes over routine tasks, employees can shift to more meaningful, value-driven roles that elevate the human contribution at work. Leaders should value human outcomes more based on creativity when they work alongside agentic AI.
- **Ethical oversight and accountability:** Leaders must ensure AI agents are used ethically, transparently and within governance frameworks.



6. What are the key considerations for ensuring the responsible deployment of Agentic AI?

Agentic AI systems hold immense potential, but inaccuracies or biases in their models can lead to error amplification, posing significant challenges. A single error in a multi-step process can propagate and magnify throughout the decision-making chain (e.g. failure to consider real-world complexities, such as cultural or situational nuances, can result in misinterpretations). As a result, it can lead to the following:

- **Ethical concerns:** Amplified biases can lead to discrimination, inequality or unfair treatment, disproportionately affecting vulnerable groups.
- **Operational failures:** Errors in critical applications, such as autonomous systems, healthcare or financial markets, can cause large-scale disruptions.
- **Erosion of trust:** Repeated inaccuracies or biased decisions damage user trust in Agentic AI systems, impeding adoption.
- **Legal and compliance risks:** Amplified errors can expose organisations to legal liabilities and regulatory sanctions under EU and national law.

Error amplification in the Agentic AI framework is a complex challenge requiring a multidisciplinary approach. By addressing the root causes of inaccuracies, employing robust mitigation strategies and aligning AI systems with European societal, legal and ethical norms, organisations can responsibly and effectively harness the potential of Agentic AI.

Responsible deployment of Agentic AI in Luxembourg and the EU requires a holistic ecosystem where policy, ethics, technology and local context converge. Responsible deployment of the Agentic AI framework should address:

- **Bias mitigation:** Promote research, guidance and awareness programs that support embedding fairness constraints and bias mitigation strategies into multi agent workflows, in line with the EU AI Act and anti discrimination rules, ensuring impartiality in decision making, especially in sensitive sectors such as healthcare and finance.
- **Human in the loop (HITL):** Ensure appropriate human oversight in the Agentic AI decision making process, particularly for high risk or critical tasks, as required by the EU AI Act. Embedding human in the loop mechanisms helps maintain accountability, minimises self reinforcing errors through corrective feedback and breaks cycles of error amplification.
- **Transparency:** Incorporate traceable decision logs, explainable reasoning, clearly defined behavioural bounds and shared accountability between stakeholders. Ensure clear labelling of AI generated content and disclosure of AI involvement in decisions, in line with the EU AI Act, DSA and sectoral guidance.

- **Privacy:** Ensure compliance with GDPR and national data protection rules, supported by privacy by design, anonymisation and pseudonymisation techniques where appropriate. Implement impact assessment and measurement frameworks (e.g., DPIAs, fundamental rights and AI risk assessments) to evaluate societal and economic outcomes.
- **Risk identification tools:** Use structured risk assessment frameworks and incident reporting mechanisms consistent with the EU AI Act, NIS2 and relevant supervisory expectations (e.g., from CNPD, CSSF and other competent authorities) to anticipate, monitor and address failures, vulnerabilities or misuse.

Furthermore, as the deployment of the Agentic AI technologies expands, ensuring compliance with existing laws becomes essential. Moreover, standardised frameworks need to be tailored to different industries/sectors. These should include:

- Regular monitoring of AI applications to ensure adherence to legal standards.
- Developing frameworks that facilitate accountability among stakeholders involved in the AI ecosystem.
- Encouraging self-regulation within industry players to complement government initiatives.
 - Auditability and explainability protocols
 - Sectoral ethics boards for oversight
 - Agile innovation models to adapt to evolving AI capabilities

Key consideration for deploying Agentic AI responsibly:

• Institutionalise oversight:

Strengthen sub-committees, create AI ethics councils and mandate regular audits using techno-legal tools.

• Empower enforceable standards:

Anchor transparency, accountability and fairness in line with the EU AI Act, the GDPR, the Digital Services Act (DSA), and relevant sectoral rules, complemented by Luxembourg's national legislation and supervisory guidance (e.g., CNPD and CSSF where applicable).

• Prioritise local context and inclusivity:

Embrace RRI principles, indigenous fairness and multilingual/localised LLMs (e.g., Sarvam AI, Dhenu) to anchor AI in cultural diversity.

• Foster multi stakeholder collaboration:

Promote industry self regulation and codes of conduct, encourage reporting and monitoring of AI related incidents, and strengthen collaboration between regulators, industry, academia and civil society. Align with European and international best practices and frameworks (e.g., EU AI Office guidance, OECD, UNESCO), ensuring they are applied in a way that reflects Luxembourg's specific economic and regulatory environment.

Conclusion

Agentic AI represents a fundamental shift in how organisations harness AI to drive meaningful outcomes. It integrates effortlessly into existing systems, tools and workflows, enhancing current operations.

They are intelligent collaborators, and the synergy between human expertise and machine intelligence can create a more agile, responsive hybrid workforce that can solve multifaceted challenges.

Organisations must approach the implementation of Agentic AI strategically to fully realise its transformative potential. Below are key recommendations to maximise its impact and return on investment.

Recommendations

- 1. Use multi-agent workflows in core business functions:** To harness the full potential of Agentic AI, embed multi-agent workflows into core business processes. These workflows act as intelligent collaborators, streamlining operations, improving responsiveness and delivering measurable value. When strategically deployed, they become key drivers of ROI and innovation.
- 2. Build reusable agents and modular agent systems:** Scalability starts with reusability. Design single agents to perform specific tasks efficiently and build a library of reusable components that can be orchestrated into larger multi-agent systems. This modular approach reduces development time, enhances maintainability and accelerates deployment across use cases.
- 3. Integrate seamlessly with your existing tech stack:** Agentic AI is designed to enhance and not replace current infrastructure. It can absorb and amplify existing automation, machine learning models and rule-based systems. Integrating with what you already have creates a unified, intelligent ecosystem that evolves with business needs.
- 4. Build long-term AI-capable architecture:** To gain a long-term competitive advantage through AI, organisations must architect systems that can seamlessly scale and integrate Agentic AI capabilities over the next five years. This requires a forward-looking approach to infrastructure, data strategy and interoperability, ensuring that tomorrow's autonomous, intelligent agents can operate across the enterprise without friction. Building AI-ready architecture is not just about technical preparedness but aligning technology roadmaps with strategic intent to future-proof the organisation's ability to innovate, differentiate and lead.
- 5. Build native Agentic AI function and capabilities:** To secure a sustainable competitive edge in an AI-driven future, organisations must build native Agentic AI capabilities. Autonomous, adaptive systems tailored to their unique data and goals will yield greater benefits than relying solely on external providers who offer the same solutions to others. Native AI ensures strategic control, exclusivity and the ability to innovate quickly, turning AI from a generic tool into a proprietary advantage. In a landscape where differentiation is everything, owning your AI stack is not just a technical choice but a strategic necessity.
- 6. Scale frameworks beyond pilots:** Build a unified lifecycle and ecosystem governance infrastructure across private and public sectors. This requires data readiness, model deployment, ethical oversight, risk management and continuous monitoring. A collaborative governance model will ensure consistency, accountability and trust, accelerating the shift from experimental pilots to responsible, mission-critical Agentic AI deployments across industries.

Contact



Cédric Jadoul

Partner, Intelligent Automation, Cross-Industry
(Financial Services, Public Sector, Private Sector)

cjadoul@deloitte.lu

+352 451 453 035

Contributors

Chintan Dharmani

Mohammed Riyaz Ahmed

Rishi Malhotra

Sagarika Mamik Gupta

Sarthak Rout

Srishti Deoras

Acknowledgements

Rahul Dhuria

Ruchira Thakur

Shreya Rajopadhye

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte provides leading professional services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets and enable clients to transform and thrive. Building on its 180+-year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 470,000 people worldwide make an impact that matters at www.deloitte.com.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.