ECONOMIST ENTERPRISE

Agents of change

Rise of the autonomous AI enterprise

Sponsored by ORACLE | KPMG

Contents

About the report

Agents of change: Rise of the autonomous AI enterprise is a whitepaper from Economist Enterprise, sponsored by Oracle and KPMG, that examines how business leaders and decision-makers are preparing for the next phase of artificial intelligence (AI) evolution, marked by the rise of agentic AI. This whitepaper is based on in-depth interviews with AI leaders to chart adoption across finance, operations, human resources (HR) and strategy, with an aim to understand best practices and pitfalls.

Economist Enterprise wishes to thank the following experts who shared their time and insights for the development of this paper:

- **Bhavesh Dayalji**, chief AI officer, S&P Global
- **Rohit Dhawan**, director of AI, Lloyd's Banking Group
- **Chandhu Nair**, senior vice president of data, AI and innovation, Lowe's
- **Kaushal Kurapati**, head of AI Agent Platform, Oracle
- **Sanjay Sehgal**, global Oracle practice leader, KPMG

This whitepaper was produced by a team of Economist Enterprise researchers:

- **Vaibhav Sahgal**, project director
- **Ailia Haider**, project manager
- **Adam Green**, writer

Executive summary

AI agents are expanding from pilots to broad deployment across enterprises. Once limited to well-defined and rule-based tasks, today's agents can reason, make decisions and interact with each other. As adoption accelerates, organisations are realising the critical importance of enterprise readiness.

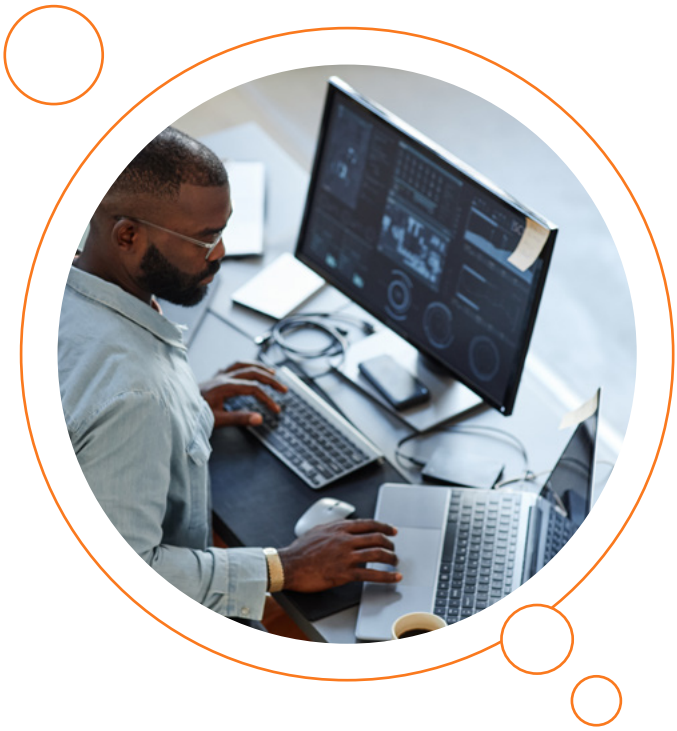

Agentic AI leaders share key characteristics. Data quality and integration are treated as strategic infrastructure. Good governance and oversight are embedded from the outset. Business objectives are well-defined and key performance metrics are measured. Lastly, they view agents as a means to support and strengthen the workforce. Taken together, these foundations enable agents to act reliably, scale safely and generate value.
To support decision-makers, this Economist Enterprise whitepaper, sponsored by Oracle and KPMG, draws on in-depth interviews with AI leaders to chart agentic AI adoption and strategy across functions, understand best practices and pitfalls, and identify the shared traits that define success. Key findings from the research include:

- **Agents provide an opportunity to redesign workflows rather than merely automate existing processes**. The most successful agentic AI implementations rebuild processes around agents, with clear hand-offs, thresholds and approval pathways, informed by the relative merits of AI versus humans across tasks.
- **Agents offer value far beyond productivity.** While raw productivity is often the first and most obvious appeal, agents offer richer commercial benefits. They can perform tasks simply not possible for humans, such as deep research reviewing thousands of sources in minutes.

Pattern recognition means they can spot anomalies that would evade the human eye, whether faulty parts in a manufacturing production line or suspect financial transactions. And they can eke out commercial advantage, like sourcing materials from more cost-effective jurisdictions or helping HR spot unconventional talent.

- **Agentic AI requires a well-governed digital estate**. Pilot projects can succeed within well-defined parameters, but to scale across functions, organisations need a well-defined governance and IT architecture that safeguards key performance requirements like observability, auditability, accountability and evaluation. Poor governance will hamper visibility and expose organisations to information security risks.
- **Flexible infrastructure and unified data help agents move freely**. Data fragmentation, inconsistent application programming interfaces (APIs) and undocumented business rules all constrain agent capabilities, especially more ambitious visions like multi-agent systems. A unified integration layer, powered by clean and consistent data communicated via standardised protocols can orchestrate workflows across systems.
- **Human capability is a critical determinant of success**. Organisations that invest in training, build centres of excellence, identify internal AI champions and define clear intervention pathways can adopt and scale agents safer and faster. Leaders look to agents as a means to augment human workers, not replace them.

Learn, adapt, act

Agents are writing the next chapter of the AI revolution. The limited chatbots and query-helpers of yesterday have evolved to be increasingly capable of reasoning, taking actions and interacting with each other. Today's agents are far more flexible and capable than traditional robotic process automation tools because they can learn, adapt and act in the real world, doing things such as accessing data and completing transactions.

Agentic AI marks a shift beyond generative AI, explains Rohit Dhawan, director of AI at Lloyds Bank. "With generative AI, you ask a question and get an answer," he says. "But an agent is code that can act like a human making decisions or finishing tasks for you."

Firms are weaving agents into daily work to cut routine tasks and lift productivity. According to Mr Dhawan, Lloyds' 50 use cases of generative AI brought in £50m of added value in 2025 alone and the organisation is now "mobilising for agentic AI". The rush to build agents is clear. A quarter of firms offering agentic AI solutions launched since 2023 have reached market faster than usual. Meanwhile, more than half of the spring 2025 cohort in Y Combinator, the legendary accelerator behind the likes of Airbnb, Dropbox and Stripe, are building agentic AI solutions.[1]

Yet adoption is not always smooth. Andrej Karpathy, an OpenAI co-founder, has warned of a glut of weak agent tools, arguing that the field is still immature.[2] As agents gain access to data and systems, cyber and information-security risks grow. A swarm of agents can also push up compute costs, add tool sprawl and unsettle staff. But when implemented correctly, the advantages are well worth the scaling challenges. This Economist Enterprise whitepaper draws on expert interviews to capture the state of play among agentic AI leaders, with a focus on three crucial functions — finance, operations and HR. It outlines the practices that help firms make agentic AI successful.

Optimised operations

Organisations are facing a tough operating environment marred by weak growth, tight margins and stubborn inflation. Many now turn to agents to cut waste and fine-tune their operations and logistics. Adopters report a 30-50% acceleration in business processes and a 20-30% improvement in workflow orchestration, not just because an agent replaces a task but because it orchestrates work that previously sat among fragmented systems.[3]

Productivity is a leading metric for operations success, alongside improved forecasting, lower cycle time, resource optimisation, reduced errors, risk mitigation and anomaly detection. Fast-moving and time-sensitive industries like retail are leading beneficiaries. "Across our business, we're evolving many of these models into agentic AI systems, capable of dynamically adjusting to new data, customer needs and operational variables in real time," says Chandhu Nair, senior vice president of data, AI and innovation at Lowe's, a US retailer.

Manufacturers already know what 'good' looks like. Stable inputs, clear data and measurable outputs let agents plan maintenance, spot faults and trim energy use. Predictive-maintenance agents now make up 38% of agentic AI in factories and have helped cut outages by 23% in some cases.[4] At Bosch's Changsha plant, specialised agents have curbed energy use by 80% and lowered emissions by 14%.[5] In pharmaceuticals, one firm cut pre-trial cycle times by 80% by sequencing pre-clinical tasks. In hospitality, automated workflow orchestration trimmed running costs by more than 10% while lifting revenue by 2-4%.[6]

Another part of this equation is procurement. Procurement teams must make quick choices in a world of shifting tariffs, supply-chain breaches, and stricter rules on sustainability and labour rights. Agents can reset stock levels, classify suppliers, review contracts, steer fulfilment and manage exceptions.[7,8] They can also improve order management, risk mitigation and product lifecycle management.[9] Lower spend, fewer manual touchpoints and a shift away from routine work are among the common measures of value. By easing the strain on supply chain teams, agents free people to focus on strategic decision-making.

Usage is growing. C.H. Robinson, a logistics firm that employs a co-ordinated group of agents to handle pricing, freight classification, order processing and tracking, has reported a 30% jump in productivity since 2023.[10] These gains come not from bolting agents onto old systems but from redesigning work around them, with clear handovers and revised approval rules.[11]

Fitter finance

Finance teams have been early adopters. They track cost per invoice, days to close and forecast accuracy, which help show whether agents earn their keep. A 2025 survey of UK financial institutions found that over half plan to increase AI investment over the next 12 months, helped by the function's clean data and structured tasks.[12] The predominance of high volume work like transaction monitoring suggest a valuable supportive role for agents.

Figure 1: Agent acceleration

Industry references to GenAI and agentic AI, 2022 and 2024

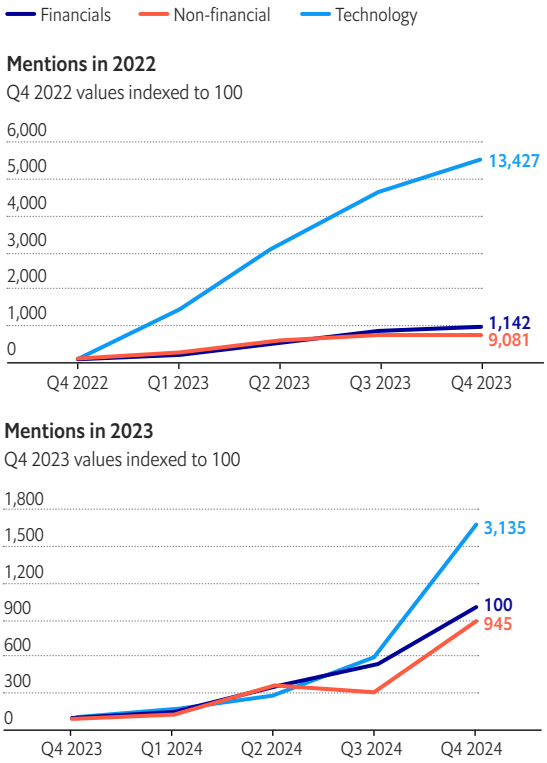

Source: AlphaSense, Citi Global Insights[13]

Notes: [1] Based on keyword searches in company documents, transcripts, and press articles for "generative AI" OR "GenAI" OR "LLM" OR "GPT" and "Agentic AI", "AI agents" and "AI bots"; [2] Non-financial sector reference above refers to all sectors excluding the financials and the technology sector.

Beyond automation, finance teams use agents for a range of knowledge-generating tasks. These include drafting investment memos, generating pitch materials, writing sector reports, strengthening scenario planning, and streamlining and integrating predictive AI with established performance management workflows.[14] A major Middle East-based bank reduced the support team's workload by 50% and operational costs by 70% after deploying a customer-facing agent, underscoring the gains from orchestration.[15]

According to Mr Dhawan, internal use cases are moving fastest in finance. "Where we have seen progress and wins are for internal use cases in back-office processes like underwriting themes, complaints and processing policy documents," he says. Consumer-facing use cases are currently limited due to regulatory restrictions preventing agents from providing financial advice.

An HR partner

HR teams are also feeling the impacts of change. Dire predictions cast agents as rivals to workers, yet many gains are shared. Agents can speed hiring, smooth employee tasks such as filing expenses and give quick guidance on handling customers. IDC, a data provider, expects that by 2026 two-fifths of roles in the largest firms will involve working alongside AI agents.[16]

Recruitment shows the shift first. A recent survey found that top uses of AI in HR include hiring, training and performance management.[17] The sheer volume of applications strains human reviewers—Goldman Sachs drew more than 300,000 applicants for 2,500 entry-level posts last year.[18] Agents can scan talent pools and review applications, raising throughput and curbing early-stage bias. They may also find strong candidates who humans might miss if they rely on shortcuts such as prestigious colleges. Metaview, an HR-tech start-up, provides AI-powered, objective interview summaries in place of messy human notes. Yet design matters: some hiring agents have reproduced the ageism and other biases seen in human screening.[19]

Agents can help after hiring too. They can flag skill gaps, suggest training and guide staff through policies on leave, IT or confidentiality. A worker can ask a bot to find a policy and then pose follow-up questions. To reap such gains, firms must build agents into processes as partners, not substitutes.

Accountable agents

Governance is the chief hurdle for firms scaling up from narrow agent tools. Many still lack a clear view of how agents should work across the business. A poll by Okta, a digital-security firm, found that only one in ten companies has a defined strategy for managing agents, even though 91% already use them.[20]

SPONSOR PERSPECTIVE

How the C-suite should lead

While many pilots are launched within business functions, scaling requires a joined-up approach—led from the C-suite. "We tend to view things within our functions, but [we] should see boundaryless organisations [that] think cross-functionally," argues Sanjay Sehgal, global Oracle practice leader at KPMG. He sees the C-suite and board, and specifically the chief financial officer (CFO), as the natural lead actor in the transition. "The C-suite, particularly the CFO, should be driving a strategic vision and overall accountability of AI in general," he explains. "Functional leads are stewards or visions of those ideas and processes."

Quality and observability top the list of concerns. "Evaluations of agent performance are still lacking," says Bhavesh Dayalji, chief AI officer at S&P Global, a financial data company. "Determining whether an agent is good or bad is a hard problem to solve." As agents act across APIs and fragmented systems, organisations lose sight of their intent and outputs. Governance must therefore sit inside the agent architecture, not be bolted on after deployment. "We don't have truly real-time, human-readable explanations of agentic AI decisions," says Mr Dhawan. That gap, he adds, weakens trust and hinders regulatory compliance in high-stakes cases.

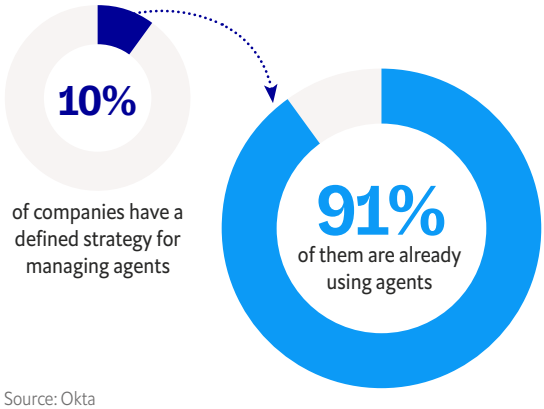

Source: Okta

SPONSOR PERSPECTIVE

The accountability imperative

Accountability ensures agents are never deployed without a clear chain of human responsibility. Kaushal Kurapati, head of AI Agent Platform at Oracle, highlights the need for "clear owners for various processes and the oversight structure", including who approves business use cases, how they are measured and what value agents deliver. He describes a three-tier oversight model, based on risk. In low-stakes processes, with high volumes, predictable workflow, a potentially high degree of manual labour and low risk, agents can act autonomously with minimal human intervention. For processes with a medium risk profile, a 'human in the loop' can choose what and when they want to review. Lastly, for high-stakes workflows, the need for high accountability means a responsible person can set key review gates and approve any final action or decision.

To solve this issue, Mr Dayalji says, organisations must grasp "communication protocols, how to connect agentic AI to third-party applications and governance within that." Agents need oversight from design through deployment. Rushing to roll them out without monitoring and feedback invites risk. Strong governance models focus on observability, auditability, accountability, traceability, testing and safeguards.[21]

In practice, observability means real-time monitoring of an agent's steps through telemetry, traces and execution logs to catch errors and 'drift'.[22] Auditability means a full record of prompts, actions, tool calls and outputs so compliance teams can reconstruct decisions after the fact.[23] This matters most in sensitive or highly regulated sectors such as finance and healthcare.

This kind of transparency lets organisations understand an agent's reasoning to ensure compliance and spot errors.[24] Evaluation checks whether agents behave safely and predictably, using benchmarks, stress tests and red-teaming, in which information security teams seek out vulnerabilities to identify gaps. These can expose edge cases and sharpen guardrails.[25]

Safeguards set hard limits on what an agent can do, regardless of its intent. Without those guardrails, cybercriminals could prise sensitive data from an agent or feed it data to influence decisions to the hacker's benefit.[26]

"Evaluations of agent performance are still lacking. Determining whether an agent is good or bad is a hard problem to solve."

Bhavesh Dayalji, chief AI officer, S&P Global

Tight access rights, activity monitoring and anomaly detection are among the protective mechanisms needed.[27] Firms need not reinvent the wheel; they can, for instance, mirror controls used for staff. "If you are a financial analyst and you can access three systems, the agent should also be able to access only those three," says Mr Dhawan.

Weak governance heightens financial, organisational and reputational risk exposure and reduces investment returns.

SPONSOR PERSPECTIVE

Plan for rain

Stress-testing is one means to improve accountability and mitigate risk but it must accommodate a broad range of future cases. Kaushal Kurapati, head of AI Agent Platform at Oracle, warns against testing only for the "happy path", thereby missing real-world variations. Potential quality issues could be identified early if agents are exposed to a wider range of scenarios without checkpoints along the way. "People often skip guardrails in pilots," he says. "Then they go into production and see that there are more possible problems and risks that can arise in task execution." Organisations tend to "over-generalise the pilot solution without looking at all of the scale, security and integration issues" that could occur, resulting in fragile systems that fail to meet user objectives. He stresses that organisations should begin by reviewing the pilot metrics and evaluating agentic AI performance against business goals, accuracy, cost and token consumption, latency, user preparedness, and scaling suitability. Subjecting agents to rapid iteration cycles and continuous monitoring to capture real-world behaviour further helps agents scale more effectively.

The agentic-ready enterprise

Accountability and safety underpin agentic AI success, yet organisations also face the large task of adapting their IT architecture. Agents can scale only if the underlying architecture enables them to retrieve data consistently, move across silos, and interact reliably with applications and APIs. This calls for a staged approach that builds the infrastructure agents need to operate safely and deliver value. The aim is a joined-up architecture that sets out the cloud environment, core applications, data layers and reporting tools.

Such upgrades are difficult to execute in large, complex organisations. "Agentic AI does present challenges from an infrastructure perspective," says Mr Dhawan, "especially when you've been through multiple technology transformations."

"Agentic AI does present challenges from an infrastructure perspective, especially when you've been through multiple technology transformations … we need to make sure the technology is actually available with our strategic vendors."

Rohit Dhawan, director of AI, Lloyds Bank

SPONSOR PERSPECTIVE

Battling architectural complexity

Agentic AI adoption is currently a work in progress, in part due to architectural complexity. "The challenge is that most companies operate in complex, multi-cloud environments with hybrid application platforms, both on-premise and cloud. They are simultaneously trying to orchestrate a broad AI strategy while developing the necessary architectural design. This makes the current state very complicated," argues Sanjay Sehgal, global Oracle practice leader at KPMG. "Organisations need a clearly defined enterprise-wide system architecture that encompasses the cloud infrastructure and underlying applications, like enterprise resource planning systems, as well as the data and reporting layers." This architecture must clearly define the roles of each application, as agentic AI capabilities are being built within existing applications and on top of these platforms by other vendors.

This might include, for instance, checking vendor readiness before deployment. "We need to make sure the technology is actually available with our strategic vendors."

Data hygiene

Data fragmentation remains a major barrier to AI adoption. Legacy systems, inconsistent data formats, undocumented business rules and complex approval chains all limit agent performance. In one survey, more than three-quarters of C-suite executives cited data integration as a top obstacle.[28,29] Disjointed and siloed data across the organisation directly impacts both deployment speed and business results. A clear and structured data audit is the place to start.

Mr Dhawan argues that agentic AI systems need "modular and distributed architectures" with "curated, cleansed and well-governed data". Even with robust data pipelines, he notes, agents can "falter when they're exposed to ambiguous user intent, incomplete data or real-world [data] distribution shifts." In practice, when data quality is inconsistent, testing must validate outputs. To build trust, firms need technology with "deep grounding [and] context awareness", which provides the ability to tie responses back to reliable data while interpreting intent and context.

Clean, standardised and well-governed data underpins accuracy. Metadata matters too. Mr Dayalji stresses the need for "the right metadata within a data fabric", which determines whether agents can fetch the context they need for deployment and evaluation. Context is critical. Without deep access to a company's operational landscape, an agent's decisions become educated guesses at best, costly mistakes at worst.

Cutting the sprawl

The next step is managing integration sprawl—controlling the rapid proliferation of disconnected AI tools, workflows and interfaces across the organisation. Too many disconnected or poorly documented API connections limit the sphere of activities that agents can perform. When APIs are inconsistent, agents struggle to access the systems they need or adapt to process changes. A mature API strategy creates a single governed layer with stable integration pathways,[30,31] reducing the need for custom one-off connections that hinder maintenance and scale. A unified platform helps create a decentralised agent mesh, enabling distributed agents to coordinate securely and handle complex tasks across enterprise resource planning, customer-relationship management, cloud and ticketing systems.[32,33]

Agent orchestration tools are still evolving, though "most current orchestration is bespoke, fragile and fails with high complexity or scale,"

explains Mr Dhawan. Some organisations run several agents, but they work in sequence or within a single interface rather than collaborating freely. Until orchestration evolves, agents, in his view, should be "restricted and bound" by tight workflows, not granted broad autonomy.

To avoid new silos, Mr Dayalji calls for "a standard set of communication protocols" for data retrieval along with a set of tools and guidelines to ensure consistent behaviour. Without the ability for IT systems, devices, or applications to communicate, exchange data, and work together, teams risk building an API system riddled with unauthorised 'shadow' or outdated 'zombie' connections that evade governance and become difficult to secure. Continuous API discovery and mandatory audits help maintain control. Enterprises can also seek operational support in these areas from vendors.[34]

From readiness to production: scaling pilots

Once data and API foundations are in place, IT and business teams should jointly target high-value, low-complexity use cases for pilots. Quick wins prove value. Clear processes and metrics guide success. Mr Nair warns against "death by a thousand use cases" that fail to deliver, using Lowe's as a positive example. The organisation "ruthlessly evaluates every AI use case on scalability and value potential. If it doesn't move the needle, it doesn't get built."

A successful pilot does not automatically translate to production. Moving from controlled environments to real workflows requires validating integrations, security requirements and workload demand before scaling.

A successful pilot does not automatically translate to production. Moving from controlled environments to real workflows requires validating integrations, security requirements and workload demand before scaling.

A hybrid workforce

As agents scale across core workflows, organisations need to prepare their workforce. Workers must learn how to supervise agents, read their behaviour, know when to escalate or override decisions and understand the limits of agent access.[35] Lowe's has invested in an AI curriculum within the Lowe's University platform, says Mr Nair. The aim is to help staff see how agents fit into daily work.

Organisations can establish a Centre of Excellence, staffed with experts to guide early rollouts. A staged approach builds competency. Mr Nair describes Lowe's "AI Transformation Office", which draws together product, engineering, legal and operations teams to ensure responsible scaling with human checks at each step.

Concerns about job security vary across the enterprise. HR and administrative teams show the greatest sensitivity. A recent attempt by Lattice, an HR platform, to add agents to its system met resistance.[36] Yet the question does not just concern direct displacement, but also how agents might change the nature of human work. One positive is the relief agents could bring to overburdened workers. In a Stanford study of 1,500 workers, nearly half favoured AI automation of their tasks, with 69% citing the desire to free up time for high-value work. Only 28% expressed concerns about agentic AI adoption, with primary barriers being a lack of trust in AI accuracy (45%), fear of job replacement (23%), and concern about preserving human judgment and creativity (16%).[37]

Lloyds assigns a team – including the CIO, COO, business owner, risk lead, HR partner and architecture lead – to each initiative. This ensures that security, risk, change and value are considered from the start.

Figure 2: Embracing agentic AI
Employees' motivations for AI agent automation*

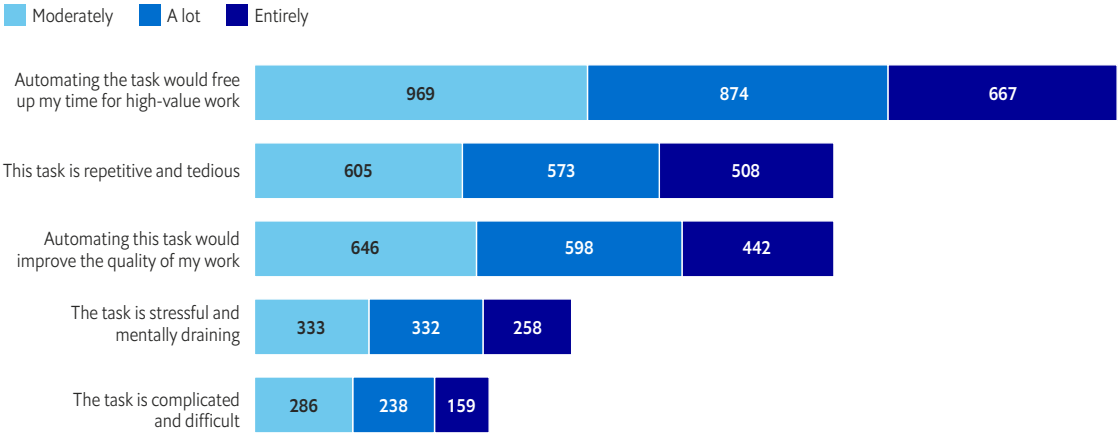

Source: Stanford University[38]

*N=3,618; Responses to the following survey question: "If an AI system can do this task for you completely, how much do you want it to do it for you? 1: Not at all; 2: Slightly; 3: Moderately; 4: A lot; 5: Entirely"

Evolving human roles

Agentic AI tools are reshaping jobs. Employees are shifting towards tasks that agents cannot handle: exercising ethical oversight, supervising workflows, checking outputs, and applying contextual judgement and reasoning. This shift is creating new 'in-demand' hybrid roles, such as forward-deployed engineers, who can operationalise models in real-world applications, embedded with the customer.[39] Agents can also raise staff capability. Mylow Companion, an agent built by Lowe's, helps store staff by retrieving product and stock information and guiding decisions. It shows how agents can support, not supplant, workers, according to Mr Nair.

Organisational design must evolve too. "My advice to anyone deploying agentic AI in a large organisation is to distribute responsibilities early," says Mr Dhawan. "Agentic AI is not just a technology; it should be treated as a team sport upfront." Lloyds assigns a team—including the chief information officer, chief operating officer, business owner, risk lead, HR partner and architecture lead—to each initiative. This ensures that security, risk, change and value are considered from the start.

"My advice to anyone deploying agentic AI in a large organisation is to distribute responsibilities early. Agentic AI is not just a technology; it should be treated as a team sport upfront."

Rohit Dhawan, director of AI, Lloyds Bank

Leadership roles must adapt too. "In the future, the role of a leader will be managing people plus agents," says Mr Dhawan. Leaders must understand the risks and "own that risk, not pass it to the development team." As agents become more autonomous, managerial responsibility will deepen.

"[AI adoption must be treated] with the same discipline and foresight as any other major business transformation."

Chandhu Nair, senior vice president of data, AI and innovation

Oversight

Employees must be able to intervene when needed. Clear pathways allow them to question and override agent choices when behaviour deviates from expectations or confidence thresholds.[40] These include human-in-the-loop approvals for sensitive tasks, periodic checks for routine work and routing uncertain cases back to people.[41] Escalation mechanisms ensure that urgent or ambiguous decisions go to human judgement. Safeguards help contain errors and define the bounds of agent autonomy.[42]

Training for supervision

Agent supervision requires new capabilities. Training for line of business managers and users should move beyond general AI literacy to afford them agent-specific oversight into issues like reasoning, consistency and modes of human intervention.[43] Scenario-based simulations help teams practice responding to failures or conflicting outputs. Some HR leaders now embed agentic AI into talent-development platforms to support employee development, framing agents as capability-builders rather than automation tools. As Mr Nair notes, agentic AI should not "replace human initiative, but amplify it".

Leading by example

Lastly, leaders must model hybrid human-agent workflows. As Mr Nair emphasises, AI adoption must be treated "with the same discipline and foresight as any other major business transformation." For the C-suite, cultivating clarity around the competence, purpose and boundaries of agentic AI systems is essential for safe and responsible scaling.[44]

Conclusion

Success in the agentic AI era will rest not only on technical progress but also on how well firms embed autonomous systems across business units and functions. The winners will treat agentic AI tools as a structural capability rather than a patchwork of point solutions. This means moving from pilots to broader orchestration supported by governance, change management and employee readiness. It will require a coordinated approach across the enterprise to share knowledge, tools and best practices.

Agentic AI-first firms share common traits. They treat data and integration as strategic assets, with systems that communicate well and provide consistent context for agents to operate across business units. They embed governance from the outset so that agents can be observed, audited and constrained. Strong governance is not a brake on progress but an enabler of scale.

These firms also recognise the human factor. They design and implement agents to work alongside people, build staff capability, and set clear pathways to observe, explain and manage agent behaviour. The shift is not from human to machine but towards a blend of human judgement and automated reasoning.

The road to becoming an agentic AI enterprise is not only a technical transition, but an architectural, organisational and cultural shift. Those investing early—and working with the right partners—on the underlying technical and organisational foundations, including integrating data, systems and tooling, embedded governance, and workforce training and collaboration, will be best positioned to scale safely and capture value.

References

1. https://www.cbinsights.com/research/report/ai-agent-bible/
2. https://www.msn.com/en-in/news/other/ai-agents-are-not-amazing-they-are-slop-says-openai-cofounder-andrej-karpathy-as-he-strongly-disagrees-with-ceo-sam-altman-on-agi-timeline/ar-AA1OTepY
3. https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms
4. https://www.mordorintelligence.com/industry-reports/agentic-artificial-intelligence-in-manufacturing-and-industrial-automation-market
5. https://www.bosch-presse.de/pressportal/de/en/how-bosch-uses-ai-in-manufacturing-260800.html
6. https://www.kearney.com/service/digital-analytics/article/the-age-of-the-agents-how-agentic-ai-offers-unprecedented-opportunities-to-reimagine-business-processes
7. https://www.gep.com/blog/technology/ai-agents-for-procurement-teams
8. https://www.weforum.org/stories/2025/09/agentic-ai-and-business-resilience/
9. https://www.alexanderthamm.com/en/blog/agentic-ai-in-supply-chain-risk-management/
10. https://www.chrobinson.com/en-us/about-us/newsroom/news/2025/ch-robinson-scales-fleet-of-ai-agents-past-30/
11. https://www.forbes.com/councils/forbestechcouncil/2025/08/26/enabling-enterprise-operations-to-thrive-in-the-era-of-agentic-ai/
12. https://www.lloydsbank.com/business/resource-centre/insight/financial-institutions-sentiment-survey.html
13. https://www.citigroup.com/rcs/citigpa/storage/public/GPS%20Report_Agentic%20AI.pdf
14. https://www.citigroup.com/rcs/citigpa/storage/public/GPS%20Report_Agentic%20AI.pdf
15. https://www.n-ix.com/agentic-ai-finance/
16. https://www.wsj.com/articles/ai-is-upending-jobs-corporate-tech-and-hr-are-teaming-up-to-figure-it-out-d6df71c9?mod=Searchresults&pos=1&page=1
17. https://www.shrm.org/topics-tools/news/technology/ai-adoption-hr-is-growing
18. https://finance.yahoo.com/news/goldman-sachs-job-golden-ticket-100000571.html
19. https://www.jdsupra.com/legalnews/eeoc-settles-first-ai-hiring-bias-6261293/
20. https://www.techradar.com/pro/security/ai-security-is-identity-security-how-okta-is-weaving-agents-into-the-security-fabric
21. https://www.iso.org/standard/42001
22. https://digitalthoughtdisruption.com/2025/07/31/agentic-ai-monitoring-observability-telemetry-analytics/
23. https://www.stratascale.com/resource/agentic-ai-governance-explainability/
24. https://www.nexastack.ai/blueprints/agentic-ai-traceability/
25. https://www.aigl.blog/agentic-ai-threats-and-mitigations/
26. https://impact.economist.com/tech-frontiers/ctrlz
27. https://bigid.com/blog/agentic-ai-guardrails/
28. https://cdn.prod.website-files.com/6130fa1501794e37c21867cf/665742a2a968a432c19a7737_MIT-report-AI-readiness-for-C-suite-leaders.pdf
29. https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms
30. https://www.raconteur.net/technology/ai-agents-performance-mangement
31. https://www.wiz.io/academy/api-governance
32. https://dev.to/webmethodman/the-agent-mesh-the-integration-renaissance-593p
33. https://www.ft.com/content/74e31d3e-4b50-43b2-9aa2-e53f41b776a8
34. https://www.apisec.ai/blog/secure-your-shadow-apis-best-practices-for-api-discovery
35. https://www.nist.gov/itl/ai-risk-management-framework
36. https://www.theverge.com/2024/7/15/24199054/lattice-digital-workers-ai
37. https://futureofwork.saltlab.stanford.edu/
38. https://arxiv.org/pdf/2506.06576
39. https://a16z.com/services-led-growth/
40. https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms
41. https://sloanreview.mit.edu/projects/the-emerging-agentic-enterprise-how-leaders-must-navigate-a-new-age-of-ai/
42. https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms
43. https://www.dataiq.global/articles/best-practices-implementing-agentic-ai/
44. https://www.weforum.org/stories/2025/07/ai-agent-economy-trust/

ECONOMIST ENTERPRISE

LONDON

The Adelphi
1-11 John Adam Street
London WC2N 6HT
United Kingdom
Tel: (44) 20 7830 7000
Email: london@economist.com

GENEVA

Rue de l'Athénée 32
1206 Geneva
Switzerland
Tel: (41) 22 566 2470
Fax: (41) 22 346 93 47
Email: geneva@economist.com

WASHINGTON DC

1920 L Street
NW Suite 500
Washington DC
20002
United States
Email: americas@economist.com

NEW YORK

900 Third Avenue
16th Floor
New York, NY 10022
United States
Tel: (1.212) 554 0600
Fax: (1.212) 586 1181/2
Email: americas@economist.com

DUBAI

Office 1301a
Aurora Tower
Dubai Media City
Dubai
Tel: (971) 4 433 4202
Fax: (971) 4 438 0224
Email: dubai@economist.com

HONG KONG

1301
12 Taikoo Wan Road
Taikoo Shing
Hong Kong
Tel: (852) 2585 3888
Fax: (852) 2802 7638
Email: asia@economist.com

SINGAPORE

8 Cross Street
#23-01 Manulife Tower
Singapore
048424
Tel: (65) 6534 5177
Fax: (65) 6534 5077
Email: asia@economist.com