pwc

Immersive Outlook

Catalysing intelligent autonomy to empower businesses

September 2025

Contents

04

‘India will emerge as a leader in both agentic AI adoption and innovation’

In an online interview with **Dr. Indranil Mitra**, **Sharad Kumar Agarwal**, Chief Digital and Information Officer, JK Tyres and Industries Ltd, reflects on agentic AI’s transformative role in manufacturing, the imperative to adapt to Gen Z work culture, and the need for clear return on investment (ROI) while investing in agentic AI.

28

05

Reimagining invoice validation and streamlining third-party risk management for a hyperscaler

In today’s fast-evolving business landscape, organisations are under mounting pressure to optimise financial operations and manage associated risks effectively. Agentic AI, with its advanced capabilities, offers transformative solutions to automate these processes and enable proactive third-party risk management. Through this case study, **Ankit Garg** demonstrates how agentic AI was applied to automate invoice validation workflows for a leading hyperscaler client – enhancing control and compliance.

32

01

Foreword

Raghav Narsalay and **Vishnupriya Sengupta**

Rise of agentic AI in the business symphony

If a business enterprise were likened to a grand orchestra, its CEO would be the master conductor guiding the interplay of complex systems across a business value chain with the finesse of a maestro interpreting a difficult score. Today, that baton is shared with a new virtuoso: the agentic AI coworker, whose tempo is hyper-fast and whose precision rivals that of the finest performers.

This is no mere evolution, it is a revolution in work rhythm.

In a recent PwC survey, 75% of senior executives said they believe AI agents will reshape the workplace more than the internet did.[1] More than a ripple, it indicates a tidal shift. Emerging value pools are now infused with intelligent autonomy, while traditional sectors are yielding to new growth domains born of both related and unrelated diversification.

Agentic AI is not just a tool – it is a protagonist, a conductor in the business symphony. It redefines how we lead, how we manage risk, how we serve customers and how we collaborate. Where software decisions once hinged on user counts, today's calculus is about how many workflows can be autonomously orchestrated and optimised.

This paradigm shift has captivated leaders across industries. Demand is surging for AI models that reason deeply and interact naturally – hallmarks of agentic intelligence. According to the survey, nearly 70% of executives report productivity gains and 57% cite cost savings, underscoring the tangible impact of these models.[2]

This edition of Immersive Outlook, '**Catalysing intelligent autonomy to empower businesses**', explores this strategic force in a landscape where agents and humans coexist – each shaping the other in a dynamic interplay of intelligence and intent. These agents strive to understand and fulfil fundamental human needs – to move, make, feed, care, build, fuel and power, govern and serve, connect and compute, and fund and insure. In effect, they accentuate value in motion.

The first article, '**Accelerating India's rise as a global powerhouse of autonomous intelligence**', charts India's ascent in the agentic AI landscape, powered by its internet-native population, mushrooming global capability centres, robust startup ecosystem, multilingual datasets, and policy support. The country currently accounts for around 40% of global AI agent deployments,[3] and hosts over 100 active agentic AI companies.[4] Its emphasis on practical implementation and real-world transformation is grounded in its use case–first philosophy – an approach that is actively reshaping industries by using agentic AI as a strategic lever to drive measurable outcomes and unlock new value streams.

The second article, '**Automation to autonomy: Redefining tomorrow's enterprise and its workforce**', offers a blueprint for scaling agentic AI by aligning goals, setting guardrails, building data foundations, and preparing the workforce for human-agent collaboration. Going beyond a technical roadmap, it's a call for strategic recalibration that harmonises innovation with risk, and agility with accountability. As AI agents gain greater autonomy to make decisions and execute tasks with increasing independence, they introduce new layers of complexity and risk. For instance, the HR function must reimagine its role to govern a blended workforce, one where digital agents and human employees co-create value, culture, and compliance.

1 PwC's AI Agent Survey

2 Ibid

3 Inc42, From Zomato To Infosys: Why India's Biggest Companies Are Betting On Agentic AI

4 Financial Express, Agentic AI startups see uptick in demand from smaller companies

The next piece, '**India will emerge as a leader in both agentic AI adoption and innovation**', is from the industry frontlines. In an online interview, Sharad Kumar Agarwal, Chief Digital and Information Officer, JK Tyres and Industries Ltd, reflects on agentic AI's transformative role in manufacturing, the imperative to adapt to Gen Z work culture, and the need for clear ROI and regulatory clarity while investing in agentic AI.

The concluding case study, '**Reimagining invoice validation and streamlining third-party risk management for a hyperscaler**', showcases how agentic AI automates complex workflows, enhancing control and compliance.

To truly unlock agentic AI's potential, businesses must go beyond bolting it onto existing systems or using it merely as a productivity booster. The competitive edge lies with early movers who deploy multi-agent systems across the enterprise, rethinking business models and workflows from the ground up.

We hope this edition will be your overture to a new era where intelligence is not just artificial, but agentic, autonomous, and aligned with human purpose – empowering you to unlock value pools, redefine customer experiences, and shape the future of work itself.

02

Accelerating India's rise as a global powerhouse of autonomous intelligence

With the rise of agentic AI, AI is evolving beyond traditional automation and moving towards attaining a higher level of autonomy. **Rajnil Mallik** and **Dr. Indranil Mitra** examine how India is uniquely positioned to lead the autonomous intelligence sector.

1. The rise of agentic AI

From machine learning (ML) to deep learning, to generative AI (GenAI) and now agentic AI – artificial intelligence has evolved rapidly in recent years. Today, agentic AI is able to reshape intelligent systems with its human-like ability to act, reason and collaborate. It is beginning to transform enterprises by reshaping workflows, driving accelerated innovation and enabling hyper-personalisation at scale. Multiple AI agents can autonomously collaborate and execute entire processes – at times across diverse ecosystems – with minimal human prompting. Agentic AI is truly reshaping enterprise software and software as a service (SaaS) business models.[1] While organisations would earlier make software purchase decisions based on the number of users, they are now evaluating solutions based on the number of workflows that can be autonomously orchestrated and optimised.

This paradigm shift is already capturing the attention of business leaders across industries as demand surges for AI models that are capable of advanced reasoning and human-like interaction – hallmarks of agentic AI.

A recent breakthrough exemplifies this shift: a newly released generative pre-trained transformer (GPT) model demonstrates sophisticated reasoning and intuitive, human-like responses. Described as being a step closer to artificial general intelligence (AGI), where AI mirrors human cognitive abilities, this model acts – amongst other things – as a powerful coding assistant that can build websites or apps from simple prompts.

The momentum doesn't stop there. Reflecting the growing integration of AI into everyday life, a leading telecom operator in India has taken a bold step, offering select users a free one-year subscription to a premium AI service. While it is common to bundle OTT services with mobile plans, this is perhaps the first time an AI assistant has been included in such offerings. It's a clear signal that both the accessibility and quality of AI systems are rapidly advancing.

Moves such as these are disruptive. By embedding AI into daily routines, they normalise the use of this technology and lower adoption barriers, especially for companies gearing up for strategic investments. Reports indicate that future-fit companies that invest in integrated agentic AI ecosystems, rather than isolated pilots for individual tasks, are poised to gain a substantial competitive edge.

1 The Economic Times, Agentic AI reshaping enterprise software pricing, business models

A case in point is an agentic AI-driven omnichannel contact centre developed by PwC US for a technology company which leverages real-time analytics, adaptive dialogue frameworks and predictive intent modelling to aid both human and virtual agents, anticipate customer needs and deliver hyper-personalised solutions. The system reduced the call time required to resolve queries by 25%, cut call transfers by up to 60% and improved customer satisfaction by 10%.[2]

A recent PwC US survey also highlights a clear shift toward agentic AI with 88% of executives planning to increase their AI-related budgets over the next year due to agentic AI and 79% reporting that AI agents are already being deployed within their organisations.[3] In India, a quarter of the companies currently using GenAI are preparing to launch agentic AI pilots or proofs of concept.[4]

PwC too has introduced its AI agent OS, a powerful orchestration framework that enables enterprises to build, integrate, and scale AI agents across diverse platforms, tools, and business functions. Designed to move companies beyond isolated pilots, it transforms fragmented AI efforts into a cohesive, enterprise-wide ecosystem, expediting its adoption.[5]

From automation to autonomy

Agentic AI is an overarching system that unifies multiple goal-oriented AI agents through an 'orchestrator' or 'guardian' agent, which alerts the sub agents based on the nature of the task at hand. While both agentic AI and GenAI leverage the power of large language models (LLMs) and ML for perception and reasoning, agentic AI goes beyond merely assisting users by taking the initiative and acting on the user's behalf. Unlike GenAI, which generates content based on prompts, agentic AI is able to perceive its environment and act and take decisions independently.

Figure 1: Key aspects of agentic AI

Autonomy
Agentic AI systems can operate independently, making decisions based on their programming, learning and environmental inputs.

Learning capability
Agentic AI systems employ ML or reinforcement learning techniques to improve their performance over time.

Interaction with environment
An agentic AI system interacts with its surroundings, perceives changes and adapts its strategies accordingly.

Goal-oriented behaviour
AI agents are designed to pursue specific objectives, optimising their actions to achieve the desired outcomes.

Workflow optimisation
AI agents enhance workflows and business processes by integrating language understanding with reasoning, planning and decision-making.

Multi-agent and system conversation
Agentic AI facilitates communication between different agents and tools (emails, code executors, search engines) to construct its own workflows.

Source: PwC analysis

2 PwC, AI agents can reimagine the future of work, your workforce and workers

3 PwC, PwC's AI Agent Survey

4 Business Standard, AI on its own: How tech is reshaping work as GenAI adoption grows in India

5 PwC, PwC launches AI agent operating system to revolutionize AI workflows for enterprises

Task-specific AI agents form the basis of a holistic agentic AI system where multiple AI agents within an agentic AI system collaborate and work as a 'team'.

Figure 2: Action arc of an agentic AI system in customer service

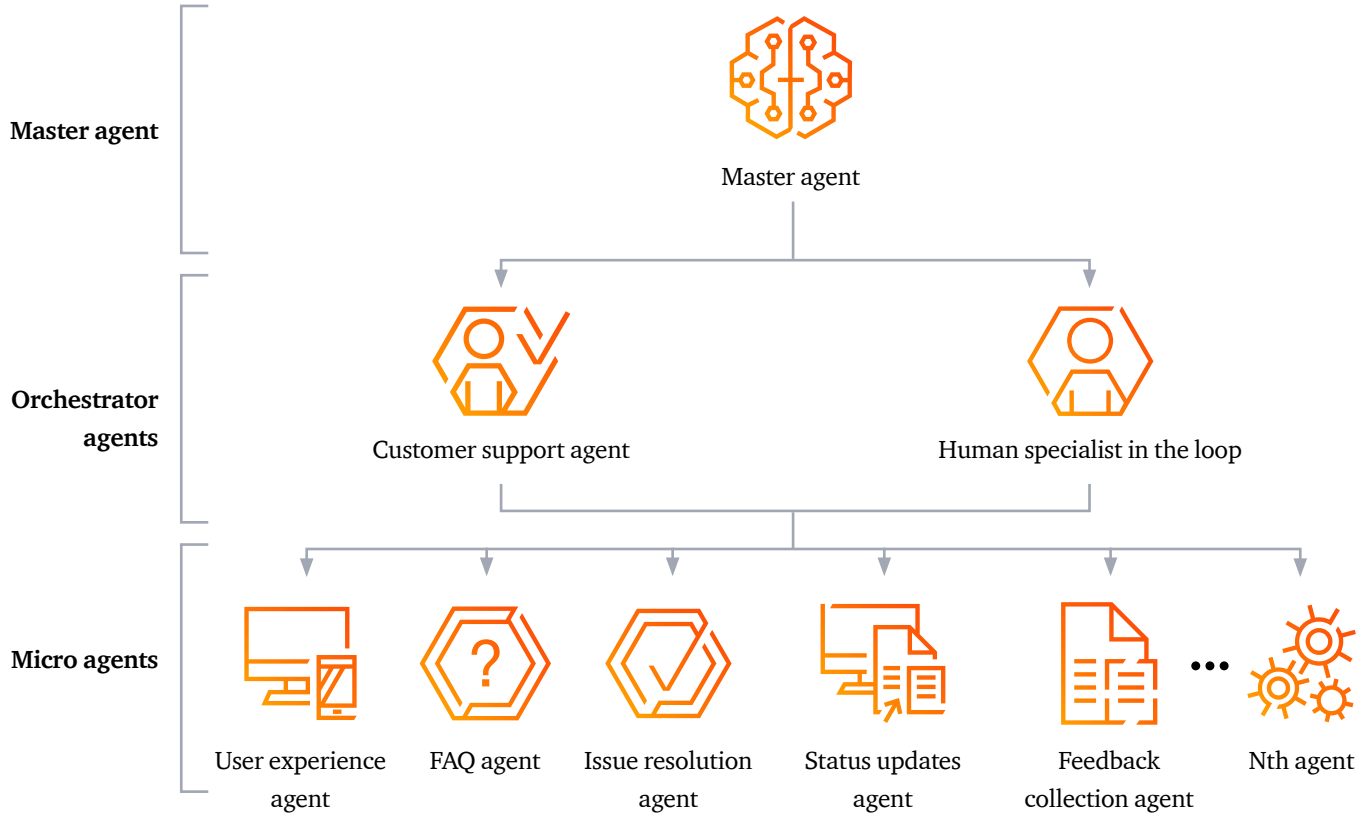

Source: PwC analysis

Agentic AI systems are proactive and take real-time decisions compared to traditional, reactive AI-based automation systems. For instance, an agent ecosystem developed by PwC US to manage a retail giant's global supply chains not only predicts inventory shortages but also negotiates supplier contracts independently.[6]

Automation systems such as robotic process automation (RPA), digital process automation (DPA) and copilots work well within rule-based setups, whereas agentic process automation (APA) can adapt dynamically to unpredictable environments and 'learn' from outcomes.

APA's impact is strong within the HR function which is gradually moving from automation to autonomy. While RPA can automate onboarding workflows by verifying documents, creating email accounts, enrolling employees in benefits programmes and carrying out payroll processing and billing, any changes to billing formats would require the system to be manually readjusted.

When applied to the same function, APA eliminates the need for manual interventions. It can autonomously initiate employee onboarding by sourcing suitable candidates from internal portals and external websites, ranking them, collecting and validating documents, and initiating background checks. Human oversight remains optional and can be exercised as needed in such scenarios.

Advanced APA-enabled tools can address employee queries, send reminders and streamline the exit process by processing final clearances and conducting exit interviews. It is important to note that the system can dynamically readjust its goals if any of the processes are altered, ensuring seamless operation.

6 PwC, AI agents can reimagine the future of work, your workforce and workers

Agentic AI, therefore, goes beyond being a mere technological evolution – it is a strategic force which is reconfiguring industries, redefining human-AI collaboration, and transforming customer experiences. The technology is based on human needs and aims to understand the nuances of human behaviour like how we move, make, feed, care, build, fuel and power, govern and serve, connect and compute, fund and insure, entailing a paradigm and value shift.

From USD 5.1 billion today, the global market size of agentic AI is projected to grow at an impressive CAGR of 44% to USD 47 billion by 2030.[7] In India, the agentic AI market is expected to touch USD 1.73 billion by 2030, surging at a staggering 54% CAGR.[8]

Uniquely positioned as a hub of innovation, cost-effective solutions and digital scale, India stands at the forefront of this intelligent transformation.

2. Agentic AI in India

India is rapidly emerging as a global powerhouse in autonomous intelligence solutions. The country currently accounts for around 40% of global AI agent deployments, underscoring its pivotal role in shaping the future of intelligent automation.[9] It also hosts over 100 active agentic AI companies.[10] IT majors in India are creating in-house agentic AI capabilities and crafting vertical-specific agentic AI solutions tailored to clients' needs. For instance, an IT services provider has partnered with an American technology company to develop AI agents across automotive, financial services, telecommunications and manufacturing industries. The agents' capabilities include content production, data engineering and workflow optimisation.

Indian enterprises are increasingly focused on enabling AI agents to collaborate seamlessly with human teams and integrate effectively into multi-agent workflows. This reflects a broader shift toward intelligent orchestration where autonomous systems complement human decision-making and operational agility. For instance, a leading IT services and consulting company intends to deploy a large group of AI agents to work alongside humans. For this, the firm aims to invest in AI data centres and cloud infrastructure and work with a human-plus-AI model. Another Indian IT giant has launched several enterprise AI agents which turn complex workflows into multi-agent business operations.

Agentic AI aligns closely with India's vision of leveraging technology for transformative impact. The country's strong IT infrastructure, dynamic startup ecosystem and proactive government policies provide a fertile ground for AI innovation. Initiatives such as the National AI Strategy and increased investment in AI research and development underscore India's commitment to global leadership in this space. India is rapidly positioning itself as a frontrunner in agentic AI, propelled by its vast pool of AI engineers and data scientists. This talent base is driving the development of autonomous systems across diverse industries, with multi-agent workflows emerging as a strategic priority for businesses seeking scalable, intelligent solutions.

3. Advantage India

India's rise in the agentic AI landscape is underpinned by the country's key strengths such as a robust pipeline of AI talent, a vibrant innovation ecosystem and a use case-first approach which prioritises practical impact. The country's emphasis on responsible AI, its growing data dividend and its status as a global hub for GCCs – many of which are at the forefront of agentic AI adoption – further reinforce its position. These advantages are supported by targeted government incentives and a regulatory environment designed to foster innovation and scale.

7 Inc42, From Zomato To Infosys: Why India's Biggest Companies Are Betting On Agentic AI

8 The Economic Times, Agentic AI finding solid traction at enterprises on clear tech use cases

9 Inc42, From Zomato To Infosys: Why India's Biggest Companies Are Betting On Agentic AI

10 Financial Express, Agentic AI startups see uptick in demand from smaller companies

a) Robust pipeline of AI talent

India is home to 16% of the world's AI talent and over nine in 10 knowledge workers in India use AI in the workplace as of 2024.[11, 12] Between 2015 and 2024, India has ranked second globally in AI skill penetration with a score of 2.5, second only to the US (2.6).[13] Moreover, the demand for agentic AI professionals in the country, driven by IT firms, startups and GCCs is expected to double to around 2,00,000 by 2026.[14]

The fastest growing job roles in India in the short-term are likely to include big data specialists, AI and ML specialists, and security management specialists. This indicates a growing demand for skills that are central to agentic AI adoption.[15] In response, Indian companies are aggressively upskilling employees in AI. Nearly nine in 10 HR leaders in India plan to reskill employees to be competitive in a market shaped by AI agents.[16] The government has also ramped up AI skilling at all levels through its flagship IndiaAI mission.[17]

Employers, therefore, must focus on nurturing India's dynamic AI talent and positioning agentic AI as a workforce multiplier and complement to human capabilities, not a replacement. With the right incentives and tools, skilled AI professionals in India can seamlessly adapt to evolving job roles and contribute meaningfully to future-ready workplaces.

b) Vibrant innovation ecosystem

India's startup ecosystem, the third largest in the world, is a key driver of the country's digital innovation landscape. The rising number of AI and GenAI startups are set to catalyse agentic AI's development and deployment. Between 2020 and 2025, over 140 native AI startups have come up in India[18] and in 2024, the number of GenAI startups in the country reached over 240, marking a 3.6x growth since 2023.[19]

The model of 'build local, scale global' is a hallmark of India's AI advantage. Initially, Indian startups offering agentic AI solutions catered to tech giants; however, now they are experiencing increased demand from small and mid-sized businesses (SMBs) as well as startups.[20] Leading Indian startups are creating low-code/no-code platforms to allow clients to build their own AI agents. One such Indian startup, an agentic AI integration platform, helps companies connect external tools to AI agents, thereby streamlining and automating complex workflows.

Some startups have been building plug-and-play AI agents tailored to specific tasks such as product discovery agents which make sense of vague user prompts and payment agents which send reminders and complete transactions. A major AI startup in India recently launched an agentic AI app which is capable of autonomously booking cabs, ordering food, paying bills and even generating in-depth research.

Thus, AI-focused startups in the country are on a steep growth trajectory. Their evolving offerings – both enterprise solutions and AI solutions for consumers – continue to capture investor interest.[21] In the coming years, Indian startups are poised to develop advanced, end-to-end agentic solutions while simultaneously transforming their own business models through increased reliance on agentic-powered digital workforces.

---

11 The Economic Times, Nine in 10 Indian knowledge workers use AI in workplace: Report

12 https://www.pib.gov.in/PressReleasePage.aspx?PRID=2108810#:~:text=India%20Ranks%201st%20in%20Global,as%20a%20major%20AI%20hub.

13 Mint, India's AI global profile is rising, but gaps remain, finds Stanford's AI Index report

14 The Economic Times, India grappling with a dearth of agentic AI experts

15 WEF, Future of Jobs Report 2025

16 ETHRWorld, Indian HR leaders expect agentic AI adoption to grow 383% by 2027: Report

17 https://indiaai.gov.in/hub/indiaai-futureskills

18 Inc42, Indian GenAI Startup Tracker: 140+ Startups Putting India On The Global AI Map

19 Financial Express, India's GenAI startups surge 3.6x in a year: NASSCOM

20 Financial Express, Agentic AI startups see uptick in demand from smaller companies

21 The Economic Times, Why are investors pouring billions into AI startups in India?

c) Use case-first approach

India's emphasis on practical implementation is reflected in its use-case-first approach to AI which prioritises solving real-world problems across sectors such as healthcare, agriculture, education and financial inclusion. This orientation is shaped by the country's vast and diverse population, which presents both unique challenges and rich opportunities for AI deployment. The impact of this approach is already visible across industries. In healthcare, agentic AI is being applied to disease prediction and personalised treatment. A case in point is a global healthcare company which deployed agentic AI workflows across its oncology practices to automate extraction and standardisation of unstructured documents, improving access to actionable clinical insights by 50%.[22]

Meanwhile, in agriculture, AI-driven solutions enhance crop yields and optimise resource use. Educational technologies powered by AI are transforming learning experiences, while AI applications in finance are expanding access to services for underserved communities. Thus, India is well-positioned to convert AI experimentation into tangible enterprise value.

d) Emphasis on responsible AI

India's focus on responsible AI and governance strengthens its position as a trusted exporter of AI solutions. The country prioritises ethical AI development to ensure that agentic systems are built with transparency and accountability. Collaborative efforts across academia, industry and government are fostering a robust ecosystem for AI research and deployment. The ISO/IEC 42001 certification provides governance and monitoring of AI-based systems within organisations. An international standard, it outlines requirements for setting up, implementing, maintaining, and constantly improving artificial intelligence management systems (AIMS) within organisations. It provides a framework for organisations to ensure that their AI systems are developed and used ethically, transparently and reliably. This standard is applicable to all types of organisations, regardless of their size or the specific AI applications they use.

Several IT majors are enabling enhanced access control measures and making AI agents auditable and policy-compliant while some are securing agent-to-agent communication to ensure safe, accountable and explainable deployment. A multi-layered guardrail deployed by an IT firm tackles bias, personally identifiable information (PII) and hallucination.[23]

Thus, Indian enterprises are increasingly prioritising AI governance, trust and adaptability with firms embedding ethical frameworks and compliance mechanisms into their agentic platforms.

e) Growing data dividend

Data plays a key role in decision-making capabilities of this autonomous technology. Lack of data readiness often stalls adoption. Given that the intelligence of AI agents is tied directly to the data they access, fragmented data landscapes could lead to ineffective decision-making.[24] A comprehensive, diverse and hygienic data lake is crucial for the functioning of agentic AI. India's vast linguistic and cultural diversity offers a powerful advantage here – rich, inclusive datasets fuel robust, multilingual AI models. Initiatives such as AI Kosh[25] under the government's IndiaAI mission and AI4Bharat – an open-source platform focused on building LLMs for Indian languages – empower developers to build context-aware, ethical AI rooted in real-world diversity. Bulk of the datasets in AI Kosh aim to support translation tools in Indian languages, such as a multilingual AI voice assistant that provides farmers with weather updates, market prices and government schemes in regional languages.[26]

22 PwC, PwC launches AI agent operating system to revolutionize AI workflows for enterprises

23 The Economic Times, Why IT needs framework for responsible agentic AI

24 ETBrandEquity, Agentic AI: The quiet revolution transforming marketing

25 https://aikosh.indiaai.gov.in/home/about-us

26 IndiaAI, AI-powered voice assistant for farmers

Establishing a robust data architecture along with rigorous evaluation processes and governance frameworks is essential for the successful implementation of agentic AI. Smooth auditing process and data sources of an agentic system must be documented in detail.[27] Extensive measures to safeguard data privacy with techniques such as encryption, relevant access permissions and frequent inspection are also important to ensure transparency.

f) Global hub for GCCs

India's GCCs are uniquely positioned to foster agentic AI innovations and scale them at a global level. Since 2020, GCCs have seen the rise of hyper automation and AI-powered automation at scale.[28] According to a recent survey conducted by PwC India, digital transformation and innovation are the leading services currently offered by GCCs to their headquarters (HQs). Furthermore, HQs are keen for their GCCs to stay focused on this area over the next five years.[29] A GCC in the banking sector deployed agentic AI to transform contact centre operations and was able to significantly cut costs and elevate customer experiences. Another GCC in the healthcare space managed to accelerate legacy modernisation processes of payer-side applications by deploying agentic AI, completing tasks that typically take years in a fraction of the time.

Thus, it's safe to say that global companies across sectors are counting on their GCCs in India to craft innovative solutions for tomorrow. GCCs in India can leverage their access to enterprise data lakes, deep domain expertise and a comprehensive understanding of their HQs' global processes to develop agentic AI solutions tailored to the parent company's specific needs.

g) Government incentives

The government's initiatives to promote self-reliance by incentivising foundational AI models, boosting the country's compute capacity and democratising access to AI through the approach of 'AI for all'[30] has set the stage for speedy agentic AI innovation.

The IndiaAI mission, with an outlay of INR 10,300 crore, aims to foster AI innovation and education by boosting compute capacity and creating a scalable cloud-computing platform for researchers and startups to train AI models. As of June 2025, over 17,000 graphics processing units (GPUs) have been installed under public-private partnerships (PPPs).[31] By offering subsidised access to compute capacity, the government aims to incentivise creation of native LLMs by startups. India is also working on building its own AI chip.[32] Under the IndiaAI mission, AI and data labs have been set up in tier 2 and tier 3 cities and AI courses have been increased at the undergraduate, master's and PhD levels.[33]

With its sustained focus on fostering an innovation ecosystem for startups, researchers and academics, India is already on the road to develop and deploy cutting-edge agentic AI solutions at scale.

4. Powering India's digital ascent for global impact

By 2028, at least 15% of routine work decisions will autonomously be made by agentic AI, up from 0% in 2024. Further, a third of enterprise software applications will include agentic AI by 2028, up from less than 1% in 2024.[34]

This indicates a fundamental shift in business processes and workflows. Over the next decade, India is poised to play a central role in enabling this global transformation.

27 PwC, Powering automation with agents

28 PwC, India's GCC story: Innovating with AI-powered automation

29 PwC, Catalysing value creation for Indian global capability centres

30 PIB, Government of India expands AI-driven skilling

31 The Economic Times, IndiaAI plan moves forward with over 17,000 GPUs successfully installed

32 Mint, India begins to design own AI chip. The tougher part comes next

33 https://indiaai.gov.in/hub/indiaai-futureskills

34 Reuters, Over 40% of agentic AI projects will be scrapped by 2027, Gartner says

While initial adoption of this emerging technology revolved around straightforward applications such as advanced customer support agents, agentic AI deployment has slowly expanded to areas with higher stakes such as healthcare management, research, cybersecurity monitoring and the more complex aspects of banking. Banks are now envisioning AI agents as lifetime engagement partners for customers – from being their first point of contact to laying out their initial financial plans to issuing loans and guiding them through investment journeys.

Our analysis reveals that agentic AI's deployment in a bank can automate KYC updates, reducing the processing time by up to 70%, simultaneously eliminating the need for paper forms, cutting cost and improving customer experience.[35] By leveraging predictive analytics, AI agents can optimise staff allocation by anticipating customer footfall and facilitate personalised customer interactions. AI agents can also significantly streamline loan processing by swiftly analysing risk profiles and validating necessary documents through APIs, reducing errors by 90% and reducing approval time by over 80%.[36]

Agentic AI is also set to catalyse a fundamental shift in India's manufacturing landscape as it can be leveraged to process repetitive manufacturing functions with precision, boosting output and reducing errors.[37]

India's proactive adoption of agentic AI will not only redefine its business and industrial landscapes but also accelerate global technological transformation, cementing the country's position as a leader in autonomous intelligence.

Contributing to this article were **Vishnupriya Sengupta**, **Ruchika Uniyal** and **Megha Adhikari**.

35 PwC analysis

36 PwC analysis

37 PwC, Agentic AI – the new frontier in GenAI

03

Automation to autonomy: Redefining tomorrow's enterprise and its workforce

As digital agents become embedded in core workflows, they are reshaping the nature of work, prompting a fundamental rethinking of governance models and decision-making structures. **Sumit Srivastav** outlines a path for strategic recalibration that harmonises innovation with risk, and agility with accountability.

1. Enterprise transformation through agentic AI

Many enterprises are being driven not by traditional automation but by agentic artificial intelligence (AI) systems that act autonomously, learn continuously and collaborate dynamically within ethical and organisational boundaries. These aren't just tools; they're intelligent partners.

Unlike conventional automation, agentic AI mirrors human capabilities such as decision-making, adaptability and contextual awareness while amplifying them with unmatched speed, precision, scalability and 24x7 availability.

Combining human creativity and empathy with AI's efficiency can result in a transformative force which can reshape leadership, build trust and drive adaptive change across enterprises.

To harness this potential, organisations must rethink how they lead, build trust and manage change. **The following sections offer a practical roadmap for deploying agentic AI through inclusive, adaptive and action-oriented strategies which can lay the foundation for a future-ready enterprise**.

Figure 1: Humans vs AI agents

Humans		AI agents
Brain Utilise experience, emotional intelligence and ethical judgement for nuanced decision-making		**Cognitive processing and decision-making** Leverage **GenAI** to analyse data, identify patterns, insights and problem-solving skills
Eyes and ears Interpret visual cues, body language and emotions to respond to environmental cues		**Information perception and analysis** Use technologies such as **computer vision** to extract and analyse information from text, documents, speech, image and videos
Speech Communicate in an empathetic way through spoken and written language for relationship building		**Interaction and engagement** Use **natural language processing (NLP)** to interact with employees and generate instant responses to queries with relevant information
Heart Use moral reasoning and empathy to make decisions, considering individual circumstances and ethical implications		**Ethical judgement and empathy** Operate within predefined **ethical guidelines** and **compliance frameworks**
Hands Physically perform tasks requiring manual intervention, adaptability and dealing with unforeseen challenges		**Task execution and action skills** Employ **RPA, workflows**, automation scripts and application programming interface (APIs) to execute tasks swiftly and accurately

Source: PwC, Powering automation with agents

A new leadership paradigm

Leading an organisation in the age of agentic AI demands a redefined leadership approach, one which cultivates trust and drives transformation by recognising AI agents as autonomous collaborators rather than mere instruments. This evolved leadership model encompasses:

- leading human–agent partnership in hybrid organisations
- establishing and sustaining trust in agentic AI
- managing organisational change through agentic enablement
- leveraging agentic AI as a catalyst for change.

Leading human-agent partnership in hybrid organisations

Agentic AI requires the restructuring of traditional leadership models into collaborative approaches where leaders guide, enable and balance human-AI offerings. This restructuring process comprises:

a) **Defining strategic boundaries:** Designing operational, ethical and regulatory peripheries is important for shaping AI behaviour and mitigating enterprise risk. By creating visible limits, organisational leaders can ensure that AI agents operate within trusted territories, thereby mitigating the risk of unfiltered autonomy while enjoying adaptability.

b) **Crafting a shared vision:** Communicating a collaborative narrative which positions AI as a provider of growth is crucial. Leaders need to frame AI deployment as a path to increase human capability, emphasising opportunities for innovation, personalisation and creative problem solving.

c) **Delegating routine decision-making:** Allowing AI to take over recurring or data-based responsibilities can enable teams to focus on high-value activities such as planning, creativity and human review. Leaders can optimise workflows by reassigning anticipated or time-sensitive functions to AI agents. For instance, organisations are using AI to tackle triaging in consumer service or manage transaction approvals in procurement. This allows humans to focus on high-value activities such as strategic analysis or client relationships.

Establishing and sustaining trust in agentic AI

Trust is dynamic in nature. It is shaped and strengthened by ongoing interaction, consistent performance and transparent communication. To inculcate trust in agentic AI, organisations must design systems which embody reliability, adapt feedback and uphold human agency. Some measures which can be taken to build trust in agentic AI are:

a) **Demonstrating competence:** Consistent and precise performance is key to building trust. AI agents are gaining users' confidence by consistently providing immense value in high-speed decision-making environments.

b) **Aligning with human intentions:** Designing AI agents to act in accordance with organisational and individual parameters enables them to act as supportive partners. AI agents which adapt to user preferences, protect data privacy and eliminate bias can strengthen trust and ethical behaviour.

c) **Increasing transparency:** Explaining decisions, unfolding algorithmic pathways and allowing human oversight can enhance users' understanding of agent conduct. By removing the 'black box' mystique, AI systems can enable collaboration rather than passive use.

d) **Practicing reciprocal learning:** Permitting AI to learn from human feedback and balance its behaviour over time is a crucial part of a collaborative and cooperative relationship between humans and AI systems. When employees provide input and AI systems tune their responses accordingly, it creates a feedback loop which encourages collaboration.

Managing organisational change through agentic enablement

Embracing agentic AI calls for a fundamental transformation – restructuring organisational frameworks, redefining roles and reinventing core processes. To ensure the successful adoption of this technology, organisations must embed adaptive change cycles supported by transparent communication, iterative design thinking and inclusive stakeholder engagement.

a) **Communicating with transparency:** Sharing goals, outlining boundaries and addressing concerns can minimise resistance and build confidence. Leaders should enable discussions around automation while focusing on augmentation over replacement.

b) **Engaging stakeholders in selection:** Involving users in selecting or testing AI tools nurtures ownership and limits pushback. When cross-functional teams are involved in pilot planning, tool assessment and success criteria definition, it leads to more authentic and accepted outcomes.

c) **Redesigning roles and upskilling continuously:** As roles evolve, the workforce moves away from routine tasks to comprehensive insights generation, and handling exceptions and unusual situations. This requires continuous upskilling of users and organisations need to be proactively involved in empowering their workforce to decipher, challenge and align their objectives with AI-generated insights.

d) **Scaling through iterative piloting:** By starting with scoped deployment, gathering feedback and adjusting the agent's operational model, organisations can minimise disruptions while increasing impact. Businesses can refine governance frameworks by aligning use cases with their strategic priorities.

Leveraging agentic AI as a catalyst for change

AI agents are more than autonomous task executors. They are change agents who are shaping new ways of working and thinking. Some of the ways in which AI agents can become catalysts for change include:

- **Identifying internal leaders:** In-depth analysis of interactive data can permit AI systems to identify employees who can act as change management agents.
- **Crafting communication:** Altering messages based on team feedback, sentiment and workflow data can increase message effectiveness and decrease disengagement.
- **Tracking engagement and growth:** Capturing key performance indicators (KPIs) such as adoption rates, sentiment scores and task fulfilment timelines allows human leaders to intervene early and guide modifications.

2. Scaling with AI agents

Though agentic AI may have moved beyond theoretical debate to becoming a technology which is being implemented at a rapid pace, its adoption is at a nascent stage. Gartner's 2025 Hype Cycle[1] places AI agents at the 'peak of inflated expectations', signalling that the enthusiasm for agentic AI may outpace practical implementations.

Figure 2: AI agents are delivering results

88%

Plan to increase AI-related budgets in the next 12 months due to agentic AI.

75%

Agree or strongly agree that AI agents will **reshape the workplace** more than the internet did.

79%

Say they are at different stages of **adoption for agentic AI**.

Source: PwC's AI Agent Survey

This combination of hype and early-stage experimentation — what some call 'pilot hell'— underscores the urgency for a structured scaling methodology. By embedding our proposed approach into their transformation roadmap, organisations can navigate the hype cycle's trough of disillusionment and achieve sustainable, enterprise-wide impact with agentic AI.

PwC India's proposed approach

Phase	Description	Key activities
Align	Tie AIagent initiatives directly to strategic objectives.	• Map high-impact use cases to business KPIs. • Secure C-suite sponsorship and budget.
Govern	Establish robust, responsible AI guardrails.	• Define ethics, security and compliance policies. • Implement risk assessments and audits.
Enable	Build technical and data foundations.	• Deploy a modular, multi-agent orchestration platform. • Create AI-ready data pipelines.
Nurture	Cultivate talent and culture for human–agent collaboration.	• Launch targeted upskilling and change management programmes. • Communicate agent roles clearly.
Transform	Pilot, measure, and continuously refine agent workflows.	• Run rapid sprints in controlled environments. • Track accuracy, ROI and user satisfaction.
Scale	Expand proven agents across the enterprise responsibly.	• Phase in autonomy levels based on reliability. • Centralise shared services and governance.

1 Gartner's 2025 Hype Cycle

The next frontier

Integrating AI agents into core operations marks the beginning – not the culmination – of transformation. As organisations move beyond enterprise-wide deployment, their AI adoption journey will evolve along three important trajectories:

1. **From single agents to swarms of collaborative agents:** Multi-agent orchestration will become the norm for tackling complex, end-to-end processes. By 2026, enterprises will deploy 'agent swarms' which can self-coordinate specialised tasks such as data gathering, decision analysis or execution by mirroring human teams but operating at machine scale and speed.
2. **Bring-your-own-AI (BYOAI) and hyper-personalisation:** Personal AI agents in employees' pockets will drive unprecedented customisation. Mirroring the bring your own device (BYOD) movement, workers will bring trusted, context-aware assistants into the enterprise, demanding seamless integration between personal and corporate environments. This fusion will unlock hyper-personalised customer experiences as agents negotiate, tailor and execute services on behalf of individuals and organisations.
3. **AI agents as strategic advisors and governance pillars:** Beyond task automation, agents will establish themselves as 'trusted advisors' offering scenario simulations, risk forecasts and strategic recommendations in real time. Concurrently, AI governance will mature into a CEO-level imperative, treating agents like APIs with rigorous auditing and security controls. Firms which master this balance of autonomy and oversight will surpass competitors, converting agent-driven insights directly into boardroom decisions.

In this emerging scenario, scaling AI agents means more than just wider deployment. It means architecting self-optimising ecosystems where agents continuously learn, collaborate and align with evolving business strategies. Organisations which embrace these trends will transform AI agents from experimental pilots into indispensable co-pilots guiding every facet of corporate value creation.

3. Sample use cases across industries

Agentic AI use cases are gaining traction across industries. While in financial services, agentic AI can act as loan processing agents that automate document verification and eligibility checks, in the retail and consumer sector they can optimise product pricing in real time by monitoring demand, competitor prices, and inventory levels. This section presents detailed use cases for agentic AI across sectors:

Banking, financial services and insurance (BFSI)

Loan processing agent

- Extract loan application documents and verify these with loan underwriting rules, automate eligibility checks to streamline approval workflows.
- Enhance process accuracy by reducing errors and efforts of credit officers while optimising the process lead time.

Claims processing agent

- Extract and segregate data from claim form, identify discrepancies to validate it against policy terms and conditions.
- Streamline the claim process and enhance fraud detection by utilising smart triaging methods.

Retail and consumer

Dynamic pricing agent

- Monitor demand patterns, competitor prices and inventory levels to optimise product pricing in real time.
- Maximise revenue and profit margins with AI-enabled pricing adjustments and promotions.

Inventory management agent

- Track stock levels and sales data across several channels in real-time and trigger replenishment orders automatically.
- Forecast demand and manage reordering point to minimise stockouts and excess inventory.

Healthcare and pharma

Query resolution agent

- Provide automated response to data-related queries during trials on the basis of protocol compliance.
- Accelerate trial timelines by reducing response lead-time and improving data-quality oversight.

Real-time patient monitoring agent

- Monitor patient vitals and movements through smart devices to alert clinicians on abnormalities.
- Provide proactive care to enable timely intervention and reduce risk for critical cases.

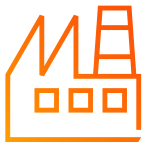

Manufacturing and industrial products

Predictive maintenance agent

- Analyse equipment data through sensor nodes and detect anomalies to warn against early signals of failure.
- Minimise downtime and maintenance costs by enabling condition-based servicing.

Demand forecasting agent

- Generate accurate, short- and long-term demand forecasting report based on historical trends and external signals.
- Enhance supply chain responsiveness by improving production planning and reducing excess inventory.

Human resources

Talent search agent

- Scan internal and external databases to identify potential candidates based on job criteria, experience and cultural fit.
- Shortlist top talent using intelligent ranking algorithms, reducing hiring cycle time and recruiter effort.

Background verification agent

- Scan employee data from education, previous employments and criminal records for automated background verification.
- Generate detailed compliance report by detecting inconsistencies for onboarding approvals.

Supply chain

Inventory monitoring agent

- Gather stock data from warehouse, suppliers and distributors to prepare an integrated real-time dashboard on inventory levels.
- Analyse and identify potential stock out and overstocking scenarios, and trigger early alerts.

Network optimisation agent

- Analyse transportation paths, delivery timelines, cost factors and recommend optimised logistic paths.
- Minimise fuel, freight and warehousing cost to improve delivery performance.

Sales and marketing

Sales development agent

- Qualify leads by analysing customer behaviour, past interactions and buying intent across channels.
- Automatically assign high-priority prospects to sales representatives and schedule timely follow-ups.

Brand monitoring agent

- Track brand mentions, sentiment and competitor activity across social media, forums and news outlets.
- Provide real-time alerts and insights to guide marketing responses and reputation management.

Agentic AI as a vendor service agent

Vendor service management has been a critical and complicated process with touchpoints across multiple departments – procurement, finance, legal and operations among others. **The process includes vendor onboarding, compliance validation, query resolution, contract lifecycle tracking and invoice/payment processing**. Most of these steps rely on manual workflows, **unstructured data exchanges** (emails, documents, spreadsheets) and **disconnected systems** (enterprise resource planning, shared folders, ticketing tools), making the process error-prone and slow.

Due to rising vendor volumes, increasing compliance mandates and cost optimisation pressures, enterprises are looking for ways to reimagine vendor lifecycle operations. The vendor service agent solution can transform unstructured processes into a scalable, intelligent and proactive ecosystem.

By orchestrating multiple specialised AI agents (for onboarding, document parsing, query handling, invoice validation and contract intelligence), this solution automates routine actions, enhances decision-making and delivers real-time insights across the vendor lifecycle.

Situations

- Significant rise in active vendor counts across industries over the past few years.
- A large chunk of vendor-related data (contracts, invoices, IDs, emails) is unstructured.
- Complex audit and compliances mandate such as Sarbanes-Oxley ACT (SOX) and General Data Protection Regulation (GDPR), and ESG-related disclosures.
- Prompt onboarding, transparent payments and consistent communication required to enhance vendor experience.

Challenges

Onboarding delays

- Vendors submit KYC, tax and banking documents via email or portal uploads.
- Manual checks for completeness and compliance often require multiple rounds.

Query overload

- Procurement and finance teams field a high volume of inquiries about payment status, PO discrepancies, or contract terms.
- Responses are inconsistent and slow when handled by different individuals.

Invoice and payment friction

- Invoices arrive as documents, scanned images, or spreadsheets and data is entered manually into ERP.
- Matching invoices against purchase orders (PO) and goods-receipt notes is error-prone and time-intensive.

Contract lifecycle blind spots

- Key dates (renewals, expirations, service level agreement milestones) sit buried in document repositories.
- Missed deadlines lead to compliance risks and unplanned cost overruns.

Solution

The vendor service agent can transform vendor lifecycle management by orchestrating a set of intelligent, task-specific AI agents which automate and streamline every step of the process by improving speed, accuracy and visibility.

The process starts with the onboarding and validation agent which extracts and validates data from vendor documents using AI-enabled document processing. It ensures completeness, compliance and direct integration into ERP systems, eliminating delays and manual rework.

After onboarding, the agent interacts with a query agent which responds to questions around payments, POs or contracts. This agent understands natural language, pulls contextual answers from enterprise systems and ensures timely, consistent communication, thereby reducing dependency on internal teams.

For invoice handling, the invoice reconciliation agent automates data extraction, matches invoices with POs and goods receipt notes (GRNs), flags mismatches and integrates the validated invoices with the financial systems. This reduces errors, accelerates payments and improves vendor trust.

A contract intelligence agent scans contracts to track renewal dates, service level agreement (SLA) terms and risk clauses to issue timely alerts and reduce compliance risks.

All agents operate under a unified orchestration layer which manages workflows, enforces SLAs and provides real-time visibility through dashboards. This centralised control ensures traceability, consistency and continuous improvement.

Together, these agents transform vendor service management into scalable, intelligent operations, freeing teams from manual tasks besides enhancing compliance and delivering a superior vendor experience.

4. Rethinking governance for agentic AI

While GenAI tools have traditionally created content, generated predictions and delivered insights based on human prompts, agentic AI introduces a new technological frontier – systems which can think and act independently. This elevated autonomy not only transforms AI agents into proactive collaborators but also brings heightened enterprise risks and governance challenges which legacy frameworks may not fully address.

In the past, governance models focused on ensuring safety, fairness and respect for human rights. Today, as organisations are increasingly adopting agentic AI to drive innovation, boost efficiency and unlock economic value, it is imperative that governance also evolves at a similar pace. New frameworks are needed to enable these intelligent systems to operate safely, ethically in alignment with human intent.

Figure 3: Similarities between onboarding and managing an agent workforce and a human workforce

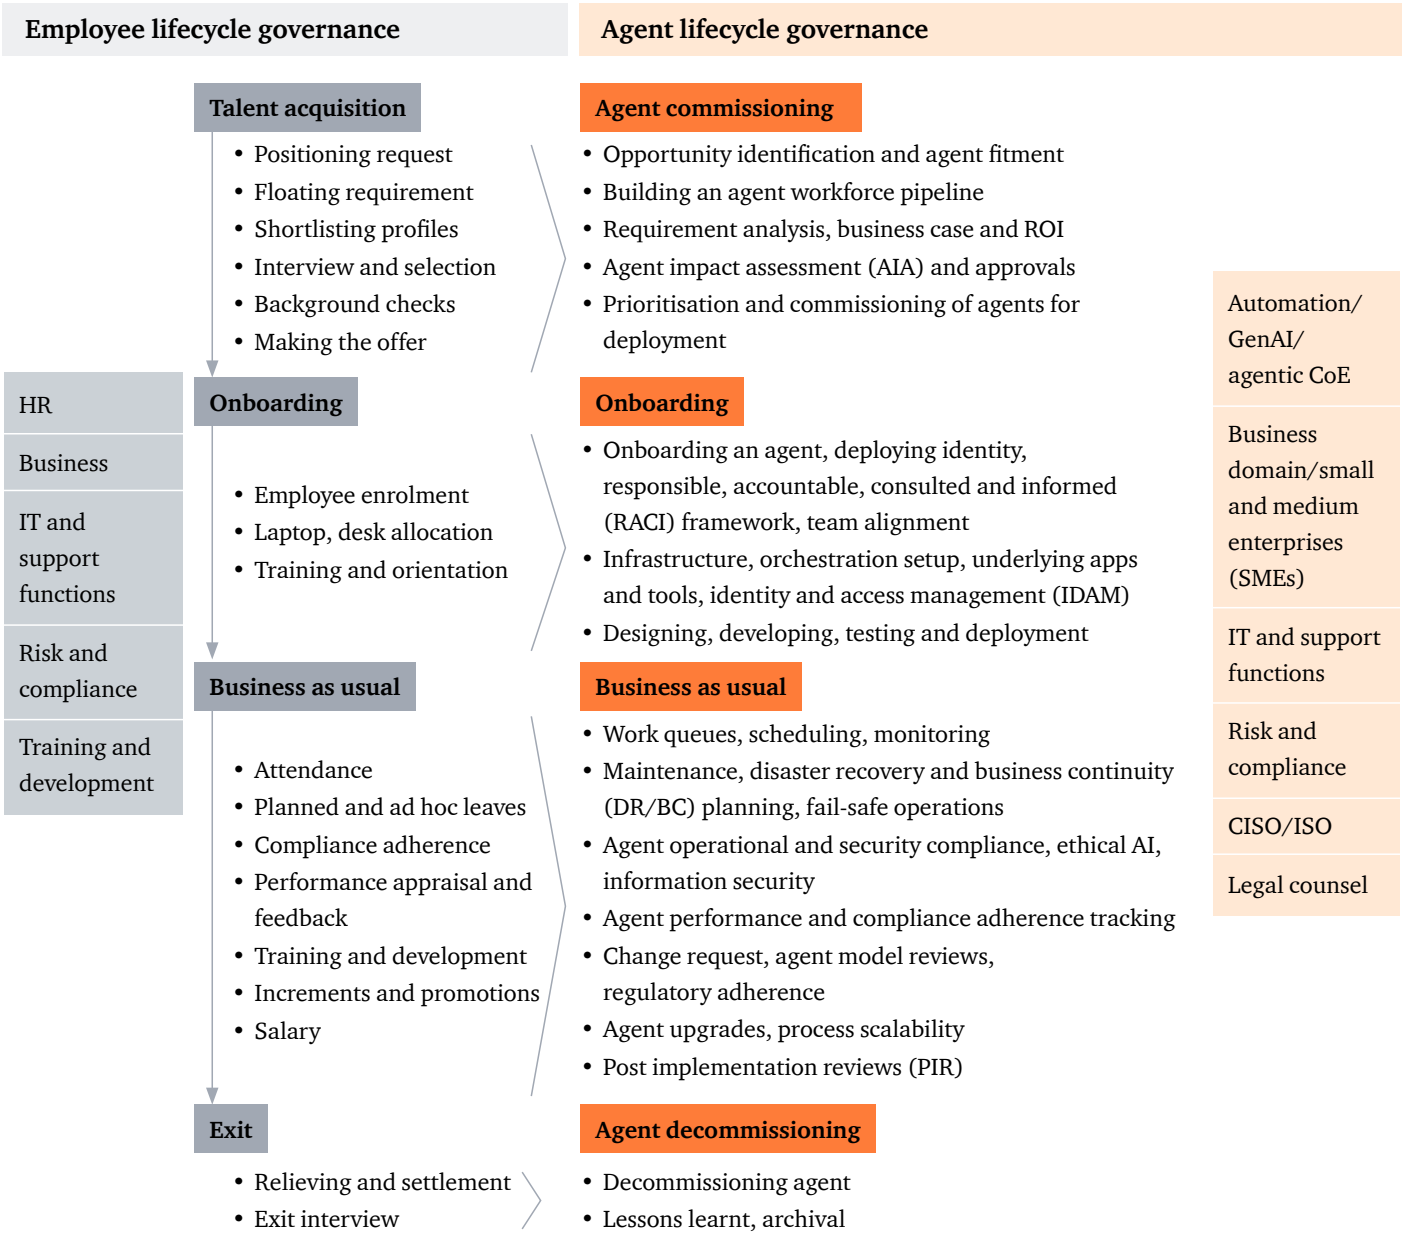

Source: PwC analysis

Integrating digital agents into the workplace goes beyond a technical implementation – it mirrors the dynamics of building and managing a team. From identifying opportunities and assessing agent performance to setting them up for success, the process is similar to the various stages of human resource management. It requires thoughtful onboarding, role alignment and continuous support to ensure that these agents contribute meaningfully to the organisational goals.

Onboarding digital agents goes beyond activating the software. It's about positioning agents with the right teams, granting them secure system accesses and ensuring that they are fully prepared to collaborate with their human colleagues. Standard operations include continuous monitoring, performance reviews, maintenance, compliance checks and system enhancements – similar to the key aspects of managing employee performance. When it comes to retiring agents and archiving their knowledge, the process is the same as followed during the traditional offboarding of employees.

Thus, establishing a digital workforce parallels the process of building a human workforce. The difference, however, arises when it comes to governing the lifecycle of AI agents. This isn't solely an IT responsibility. Instead, it demands a cross-functional approach involving automation experts, CoEs, cybersecurity specialists, legal advisors, risk and compliance teams, and domain experts. Together, these stakeholders must define the rules of engagement – how agents are onboarded, evaluated, upgraded and retired – while ensuring that the operations remain safe, ethical and aligned with organisational values. This governance framework transforms a collection of standalone tools into a cohesive, intelligent and future-ready digital team, one which empowers business functions and enhances agility in a constantly evolving business landscape.

Autonomy versus accountability

The inherent autonomy which defines AI agents also adds to the complexity of their governance. Unlike traditional software systems which followed strict rule-based programming, AI agents are designed using ML algorithms to analyse tasks and determine the necessary action based on their reasoning. This level of autonomy allows agents to function in real-time, which introduces a unique set of risks. Some of these risks include:

a) **Lack of transparency:** Reduced or lack of human inclusion in the loop makes it difficult to ensure that AI agents are acting in a fair and ethical manner. In high-value, high-risk areas, such level of independence could lead to undesired outcomes. This presents a challenge for leaders to find the right balance between the level of autonomy needed for agents and the need to control them with the right guardrails.

b) **Unintentional consequences:** As more and more agentic systems evolve with higher levels of sophistication, the decisions they make may not be easy for humans to interpret. Simpler, rule-based systems with traceable logic have a certain amount of predictability. However, AI-based decisions follow reasoning powered by complex ML models and without a human in the loop they can pose challenges such as lack of transparency, accountability, and issues related to bias and fairness. Suppose self-driving vehicles were to make poor judgements resulting in an increase in road safety incidents or a healthcare agent were to provide incorrect diagnoses. The impact could be significant, hence agentic governance frameworks become critical to make the decision-making more transparent, accountable and aligned with organisational and regulatory policies.

c) **Bias and ethical concerns:** Since historical data forms the backbone for AI systems to learn from, it is important for the data to be rid of any biases for the AI system to be free from a skewed view. Sometimes, AI systems may take decisions that prioritise goals such as efficiency, potentially at the expense of human qualities like empathy, which can be problematic in sensitive situations.

d) **Systems and data vulnerability:** Access to various data types, tools and systems also make agents more vulnerable to security risks such as memory manipulation, making them potential targets, resulting in cascading effects in a multi-agent system. These risks increase the chances of system breaches as compared to traditional AI.

Agentic AI governance framework

Governance practices which apply to traditional AI governance such as data governance, continuous risk evaluations, transparency in workflows, compliance and user awareness are also applicable to agentic AI systems. However, with agentic AI, governance frameworks need to be more advanced given the increased autonomy and sophistication of these systems.

Organisations find it challenging to implement agentic governance frameworks which are balanced and have the right level of human oversight, automation and AI-driven self-regulation. Some key considerations while deploying governance frameworks include developing frameworks with in-built mechanisms for:

- **explainability and interpretability** to ensure decisions taken by AI agents are transparent at all times and the reasoning is clear to the developers
- **bias and fairness management methods** to detect and mitigate unfair outcomes faster
- **anomaly detection and self-correction** to enable AI to autonomously correct errors or alert humans for corrective actions. Peer-to-peer agent monitoring is another way to rectify errors when issues arise.

Agentic governance should also incorporate self-learning mechanisms which continuously refine and update governance models based on user sentiment and feedback, incident response and audit reports. This method of monitoring through a continuous feedback loop allows better tracking and evaluation of agents' performance. Other suggested practices for better agentic governance being explored include:

- incentivising beneficial use of agents
- developing mechanisms and structures for managing risks across the agent lifecycle, including technical, legal and policy-based interventions.

Organisations can enhance their existing automation governance frameworks to effectively oversee digital agents and intelligent systems. To ensure comprehensive and responsible integration, governance should span seven critical dimensions:

- process governance
- data privacy and security
- monitoring and reporting
- human autonomy
- accountability and compliance
- fairness and bias
- transparency.

By focusing on these categories, organisations can create a unified governance structure which supports the safe, ethical and scalable deployment of agentic AI.

Figure 4: Simplifying agentic process automation governance

Source: PwC analysis

Preparing for the future of work

Organisations are witnessing a profound shift from human-led, manual processes to AI-powered systems where autonomous agents operate alongside people with enhanced speed and precision. This evolution promises not only cost efficiencies but also the creation of new revenue streams and scalable service delivery models.

However, as AI agents gain greater autonomy in decision-making and execution, they introduce new layers of risk and complexity. For instance, the HR function must evolve to manage a blended workforce of humans and digital agents. This will require rethinking talent strategies, developing new skill sets, and redefining how human potential is sourced, nurtured and measured. As agents take over routine and repetitive tasks, human roles will need to shift toward high-value, strategic contributions.

To navigate this transformation, organisations must strike a balance between innovation, investment and return on value. This includes developing both quantitative and qualitative methods to assess the performance of human-agent teams. Governance frameworks should also be expanded to address emerging organisational and societal risks. Enabling continuous innovation in this new landscape will require leaders to establish a robust and responsible AI framework, one which ensures safety, ethics and long-term sustainability.

Contributing to this article were **Hariprasad Gajapathy** and **Disha Thadani.**

04

‘India will emerge as a leader in both agentic AI adoption and innovation’

For **Sharad Kumar Agarwal**, Chief Digital and Information Officer, JK Tyre & Industries Ltd, agentic AI plays the role of a symphony conductor. In an online conversation with **Dr. Indranil Mitra**, Partner, iDAC (intelligent data, agents and cloud) at PwC India, he shared how agentic AI is reshaping manufacturing through process automation, workforce adaptation, and ethical deployment, underscoring the importance of trustworthy data sources and responsible AI governance.

Excerpts from the interview

Dr. Indranil Mitra

Dr. Indranil Mitra: JK Tyre and Industries Ltd. has been a pioneer in driving emerging technologies, from being the first to launch smart tyres in India to deploying AI-driven processes for agile manufacturing. How do you envision agentic AI, a technology that works with minimal human intervention, reshaping manufacturing processes?

Sharad Kumar Agarwal

Sharad Kumar Agarwal: Manufacturing has long been a neglected area in terms of digital innovations. While leveraging automation and robotics has been the forte of the manufacturing sector, using digital tools – specifically AI – in manufacturing has taken a backseat. We found that there is significant value to be gained [from deploying AI in manufacturing], which can be used to enhance the bottom line. It [AI] can help us improve productivity and sustainability, and can be factored in anywhere.

Agentic AI would play the role of an orchestrator of multiple AI agents. They can be LLM agents or RPA agents. Think of it as the orchestrator of a symphony. There are several repetitive processes in manufacturing [that can be automated]. Several robots may be functioning on the shop floor, but they need to work together under orchestration.

I also foresee a lot of industries in India moving towards dark factories. Agentic AI will help a lot there.

Dr. Indranil Mitra: How is India positioned to emerge as a leader in global agentic AI innovation?

Sharad Kumar Agarwal: Several reports say that globally, India has been at the peak of AI innovation. This industrial revolution may have started from the Western world, but India has found a sweet spot in terms of adoption. This is due to multiple reasons – the most important one being the availability of talent. The excellent talent pool that we have here [in India] is open to learning and exploring new ideas. My personal view is that business leaders are ready to experiment, and these experiments are not just proofs of concept (PoCs), they are happening section-wise.

Leaders are trying to bring in the benefits of AI, specifically agentic AI and other digital frameworks, to improve the top line, unlock cost savings and enhance the bottom line. When it comes to manufacturing, the two most important agendas are safety and sustainability. Safety is a top priority for us. As for sustainability, our company as well as our industry takes it very seriously. I'm very sure that India would emerge as a leader in agentic AI adoption. We have been followers in innovations so far, but we can emerge as the leaders.

"Leaders are trying to bring in the benefits of AI, specifically agentic AI and other digital frameworks, to improve the top line, unlock cost savings and enhance the bottom line."

Dr. Indranil Mitra: I think your point about sustainability and safety is a game changer, if AI is deployed properly. You spoke about talent. I would like to understand your viewpoint on how agentic AI is reshaping the workforce in terms of human–AI collaboration on the shop floor and how you are upgrading your talent pool.

Sharad Kumar Agarwal: We need to do away with the misconception that AI is here to replace humans. It is meant to enable humans, so humans can focus on using their grey matter. Talent is available. Today's workforce is quite young – most would be in their early or late twenties, and they are doing a wonderful job. We need to redesign our human capital management (HCM) policies to adapt to Gen Z work culture. They won't adapt [to our ways]; we should adapt together, earn the fruits of the labour and help businesses to grow. We have designed offices keeping in mind the talent that we are acquiring now or that will be coming in tomorrow. Companies must thus adapt to the talent – not the other way around.

Dr. Indranil Mitra: JK Tyre has been tirelessly modernising its manufacturing processes. You have leveraged ML and industrial internet of things (IIoT). You deployed digital twins at a time when organisations had only started talking about it. Now you have deployed AI agents as well. This has boosted your competitive edge considerably. Do you plan to expand your agentic AI deployment? Any promising use cases you can talk about?

Sharad Kumar Agarwal: We are surely going to deploy agentic AI on top of the smart factories that we have already built and are building across multiple locations. We're trying to make agentic AI a part of our virtual twins and create a combination of digital and AI twins.

We are also trying to make agentic AI systems run factories for us by enabling them to take decisions wherever human intervention can be minimised. Humans would of course be in the loop wherever the exceptions are beyond the current understanding of the AI agent. A use case that we're building is a failure mode and effects analysis (FMEA) agent. We want humans to focus on how to improve and innovate.

"A use case that we're building is a failure mode and effects analysis (FMEA) agent. We want humans to focus on how to improve and innovate."

Dr. Indranil Mitra: Agentic AI and LLMs are very powerful data analysers – capable of deriving actionable insights from large amounts of unstructured data. Do you think agentic AI can be leveraged to pinpoint specific customer preferences? With sales channels and retailers across the country, do you think JK Tyre can leverage agentic AI to enable hyper-personalisation of products and appeal to various distinct market segments?

Sharad Kumar Agarwal: Theoretically, hyper-personalisation is possible using agentic AI. However, agentic AI is dependent on the data it is trained on, so there are possibilities of data bias and hallucination. Agentic AI may not be the perfect tool [for hyper-personalisation at the moment] but it can slowly graduate and be trained [to get there]. It can't be a one-size-fits-all solution.

> “Agentic AI may enable hyper-personalisation for one product but not for another. So, it needs to be trained and retrained with sufficient human intervention.”

Agentic AI may enable hyper-personalisation for one product but not for another. So, it needs to be trained and retrained with sufficient human intervention. Humans must remove any bias they spot. It can't be purely unsupervised learning. Having said that, yes, they can enable hyper-personalisation but can only achieve about 80% of the work. Humans will need to do the remaining 20%.

Dr. Indranil Mitra: Agentic AI can often be difficult to implement and companies may struggle to get returns on what they have invested in the technology. Do you think that the hype around agentic AI exceeds its tangible value? Companies at times push ahead with implementation even when PoCs don't yield the correct results. What do you think is the initial roadblock in the enterprise-wide adoption of agentic AI and how can we remove those to enable successful implementation?

Sharad Kumar Agarwal: The base roadblock is the unavailability of data. All data streams must come together to form a data lake. For doing a pilot, we need to have a very clear-cut value on the table – detailing the investments and the ROI. We must push out the use cases that don't bring value and focus on the ones that do.

And, as I previously said, we need to segregate the foundation layer – where data availability is important. There's a cost to building the data pipeline. You may need to deploy sensors and pull data out of programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems and mobile value-added services (MVAs). You will have to integrate IP security systems.

So the foundation layer has to be thought of as a hygiene function, and the ROI has to be calculated on top of specific use cases.

Another point is sometimes people get attached to solutions, especially if they are part of the team that has built the solution. It is important to focus on the problem. The solution can be anyone's. If the problem has to be solved, and there are ten solutions, one should not be biased about one's own solution.

Dr. Indranil Mitra: You spoke about biases and ethical concerns around agentic AI deployment. Do you have any concerns about agentic AI deployment, especially when it comes to balancing the autonomy of such agents with humans in the loop? Have you placed any guardrails around agentic AI deployments to ensure safety and transparency of the systems in your organisation?

> “If an AI agent has to repair a machine, it should be fed that machine's installation and repair manual, so that the agent won't go to open-source information to find a solution.”

Sharad Kumar Agarwal: Firstly, in the manufacturing sector, we train our AI agents only on reliable data sources. If an AI agent has to repair a machine, it should be fed that machine's installation and repair manual, so that the agent won't go to open-source information to find a solution. So we're trying to narrow the focus of AI agents. Secondly, balancing agentic autonomy with the human interface varies in each case.

We need to see in which cases AI agents can be given full autonomy and where human approval is required. AI agents can be allowed 100% autonomy in some cases, but not all.

You have to see the source of the training data; it cannot be a black box. You will need to remove the biases, though it's easier said than done. It is also important to clearly look into the different AI regulations of each country.

I read that legal frameworks follow innovation. As innovation moves forward, it finds its way. There may be some bad spots – a human may misuse their comprehension capabilities and similarly, an AI agent may get misdirected. It is our job to keep them within the guardrails.

Dr. Indranil Mitra: Thank you so much for your time.

Sharad Kumar Agarwal: Thank you, Dr. Mitra. It was my pleasure. I enjoyed the conversation.

05

Reimagining invoice validation and streamlining third-party risk management for a hyperscaler

In today's fast-evolving business landscape, organisations are under mounting pressure to optimise financial operations and manage associated risks effectively. Agentic AI, with its advanced capabilities, offers transformative solutions to automate these processes and enable proactive third-party risk management. Through this case study, **Ankit Garg** demonstrates how agentic AI was applied to automate invoice validation workflows for a leading hyperscaler client - enhancing control and compliance. The study also outlines the key benefits and practical challenges encountered during the implementation of the agentic AI solution.

PwC India collaborated with a global hyperscaler to transform the invoice validation process as part of the accounts payable cycle by automating validation checks with agentic AI. The client, that operates across multiple geographies, faced significant challenges in processing thousands of invoices due to its vast and diverse vendor network. Delays in invoice processing, compliance concerns and human errors due to manual processing were identified as key risks. Moreover, conventional systems struggled to keep up with the complexities of their invoice processing operations, especially when working with non-standard invoice formats, varying demands of individual vendors and data files. Lack of consistent control mechanisms across the process was another major challenge.

The client had considered a few alternatives before adopting agentic AI; however, each of them faced the following challenges:

- **Optical character recognition (OCR) and rules-based processing:** This was effective for regular or standard invoices but challenges surfaced while handling almost 40% of the non-standard invoice formats. Varying invoice formats from multiple vendors made it difficult to automate the process to an acceptable level as it became a recurring job and exposed the organisation to huge risks.
- **ML-based solutions:** The solution helped increase the accuracy of the data extraction process but struggled when it came to contextual responses for different styles of invoices or credit notes. Reading and understanding the data and adjusting the context of modified names was a persistent concern.
- **Workflow automation:** Basic automation streamlined the routing of the workflow; however, the client was unable to improve validation accuracy and decision-making.

By deploying agentic AI along with a master orchestration framework, PwC India was able to design a bespoke, scalable automation plan to enhance controls, governance and risk management.

Implementation process

Agentic AI was plugged into the client's invoice processing and validation workflow to enable seamless extraction of data and to automate validation and compliance checks. Key functions of the agentic AI solution included:

- **Data integration:** Consolidation of structured and unstructured data from scanned invoices and multiple complex files.
- **Customised approach:** Using the client's proprietary LLM model with a retrieval-augmented generation (RAG) framework and computational engine helped process a variety of invoice formats by applying complex arithmetic and analytical operations.
- **Compliance mechanisms:** By automating checks, the agentic solution enhanced the compliance of the invoice process with the company's policies. It also assured conformance with multiple vendor contracts to effectively manage third-party risk.
- **Atomic operations and templatised prompts:** Invoice validation logic was broken down into atomic operations such as individual arithmetic operations and rule-based checks where each check was mapped to a templatised prompt. These prompts used structured metadata rather than raw data, enabling consistent and reusable logic across vendors for easier debugging and version control.
- **Agent orchestration:** An agent orchestration layer handled the flow of tasks intelligently and dispatched parts of the invoice to specialised agents depending on context and validation requirements. It enabled inter-validations to be executed in a dependent and sequential order. This was based on their creation sequence with built-in support for fall back and retry mechanisms to tackle failures and anomalies.
- **Modular architecture:** Our solution followed a modular architecture where each component — data ingestion, validation, compliance and reporting — functioned independently and allowed plug-and-play support, parallel processing, fast experiments with new prompts and LLMs or logic changes without needing system-wide changes.
- **Standardised input file templates:** The solution replaced fragile spreadsheet-based macros and traditional automation with scalable, standard templates for better consistency.

Challenges

There were several challenges in the implementation of the solution.

1. **Data quality and integration complexity:** The use of agentic AI solutions with existing infrastructure can be challenging since the infrastructure is often traditional and non-interoperable. These environments are often heavily customised for supporting cutting-edge AI, making them difficult to implement with other robots or projects and lacking sufficient computational power.

   The client's process was manual and unstandardised with discrepancies in data, misaligned columns and unsegregated data structures. These became an obstacle for generating stable templates for automatic vendor verification.

2. **LLM changes and integration:** Since agentic AI technology is constantly evolving, regular updates to LLMs may result in integration difficulties. There may also be processing capability changes with the new versions; therefore, it is necessary to calibrate and train the model repeatedly so that it is compatible with newer versions.
3. **Ethical and compliance concerns:** Responsible AI is a critical component of AI adoption. Systems must adhere to certain ethical principles and relevant laws and standards. For instance, key priorities for this project were transparency, accountability and explainability in AI decision-making which are vital factors in minimising compliance risks and fostering stakeholder trust.
4. **Template maintenance and retraining:** AI-driven systems must be retrained periodically to keep abreast of new templates, business rules and changes in vendor formats.

RAG limitations and LLM accuracy challenges

The successful use of RAG systems depends on the quality of knowledge the model can use. Misdated or inconsistently labelled data can misguide AI outputs. As a result, some of the output may be incorrect. Therefore, ongoing validation and fine-tuning are necessary to maintain high accuracy and minimise errors in production.

For instance, the LLM used in this project didn't meet some of the validation checks that it met when it processed a data file from the same vendor in a prior iteration.

A retroactive root cause analysis was completed highlighting that the validation template was included along with the data input, causing the LLM to generate a code using its computational capabilities. This led to inconsistent results and ensuring reliable validation was a challenge.

To avoid this, a governance mechanism to manage or control the code produced by the computational engine is required to maintain stable validation of results.

Implementation teams should include the stakeholders early on to decide and design data validation and compliance rules. They should also focus on training teams on AI confidence and on the fallibility to determine which cases require human intervention and review.

The learning curve

The implementation journey of agentic AI came with significant learnings:

01 **Initial over-flagging:** With cautious thresholds and a biased training set, almost 60% of the invoices were flagged in the first two weeks.

Solution: A centralised repository was implemented to maintain a uniform gateway to these vendor-specific validation files.

Key takeaway: AI models require exposure to unstructured, real-world data including edge cases to improve generalisation and reduce false positives.

02 **Vendor-specific patterns matter:** Though accuracy increased to 85% by the end of week three, there was an issue with processing vendor data consistently.

Solution: A standardised, templatised process of prompts with procedural steps to support consistent validations and avoid duplication helped mitigate this.

Key takeaway: In addition to patterns, the AI model also needs to learn vendor context to be reliable.

03 **Managing computational engine variability:** The LLM-generated code executed at runtime led to unpredicted validation outcomes.

Solution: A constrained execution model based only on LLM-generated queries with metadata-driven prompts and a regulatory layer for stability addressed this concern.

Key takeaway: Organisations must refrain from executing LLM-generated codes directly and leverage metadata-driven triggering and controlled environments for consistent results.

04 **Confidence scoring:** One of the key steps of the implementation process was to add confidence levels to all validation output.

Solution: This approach enabled greater focus on cases with high uncertainty both in terms of insights and the machine-generated output while reducing the human workload.

Key takeaway: Confidence scoring can be easily automated at scale. It also allows human intervention.

Focus points

Implementing agentic AI proved to be challenging due to data inconsistencies, evolving AI models and strict compliance expectations. Addressing these challenges required careful governance, retraining and adaptive design. Some of the factors which hindered the adoption of this technology were:

- **Challenges in achieving end-to-end automation:** The best level of automation hit a ceiling at about 95%. Full automation was impossible due to complicated cases and the level of human judgement needed.
- **Lack of change management:** Though technical deployment had its share of challenges, one of the major hurdles was convincing business users to adapt the new solution.
- **Generic rules:** A one-size-fits-all solution is not valid for inter-department and varied vendor requirements.
- **Uniform vendor treatment:** Model training and standardised templates across vendors were not possible due to the difference in vendor size, scope and customised checks.

Future areas of implementation

Besides invoice validation, some of the areas where the solution can be implemented are:

- **Predictive vendor management:** By processing historical and current data using AI, businesses can predict which vendors are likely to submit incorrect or non-compliant invoices, allowing action to be taken in advance to proactively reduce financial exposure and strengthen relationships with vendors.
- **Dynamic approval of workflows:** Future AI systems will not only be contextually aware but will also utilise dynamic approval hierarchies which can be changed as per real-time business requirements and invoice conditions while still remaining compliant.
- **Cross-process intelligence:** Information obtained from invoice validation can be used in other areas such as purchasing, contract negotiation and vendor performance management which could enable businesses to develop an interconnected IT financial ecosystem.
- **Chat-based interface:** In future, the proliferation of conversational agents will enable finance teams, risk teams and business stakeholders to ask natural language questions about invoice data, risk assessments or compliance statuses to gain access to critical data and make instant, data-driven decisions.
- **New vendor onboarding:** Smartly orchestrated agents and automated workflows could reduce time and cost of new vendor onboarding by automatically creating custom validation templates, prompt configurations and system setups, thereby reducing manual effort and accelerating time-to-value.

Steps for successful implementation

To successfully adopt agentic AI, organisations should:

1. **Start with data readiness** by providing clean and normalised data. This phase would detect and respond to structural inconsistencies such as malformed columns and header mismatch prior to implementation.
2. **Provide centralised access** by creating an easily accessible single source of truth for both structured and unstructured invoice data to track and maintain compliance.
3. **Use metadata over raw data** since raw invoice data in LLMs often yields inconsistent results. Switching to metadata-based prompts ({vendor_type}, {invoice_amount}, {currency}) can deliver consistent output.
4. **Develop a robust design and modular architecture** for the solution with the following steps:
    - **Atomic operation layer:** Break and templatise invoice validations into individual steps based on distinct business rules or calculations.
    - **Agent orchestration:** Route subtasks such as amount checks, vendor authentication or tax validation to specialised agents using a task orchestration layer which considers the context of the request.
    - **Controlled code execution:** Introduce a sandbox for code interpretation like logic or templates that are validated, restricting LLMs from generating and executing dynamic codes at validations.
5. **Templatise prompts** for scalability and faster debug by:
    - **Using prompt templates per vendor/use case:** Leveraging generic prompt templates is effective when applying identical validation checks for multiple vendors. This cuts down the need for re-writing of the check validation template, thus streamlining the process and adding robustness to the system.
    - **Maintaining prompt versioning:** Tracking the evolution of prompts over time, including rollbacks and experimenting, should be part of the implementation process.
    - **Scaling by segment:** Starting with vendors that have similar formats or validation rules before rolling to difficult or high variance sources is critical.
6. **Plan template drift** and model variability by:
    - **Monitoring LLM output variability:** LLMs can produce different answers even when they are given the same input across sessions or releases. Therefore, it is important to implement fallback checks based on deterministic logic to ensure consistency and reliability regardless of changes in LLM behaviour.
    - **Implementing lightweight retraining loops:** Keeping track of validation fails and incorporating re-try in feedback loops with a set limit on the number of retry attempts is essential.
    - **Introducing confidence scoring:** A confidence level should be attached to every validation outcome. This enables the user to scale people-powered automation and detect low-confidence anomalies for reconsideration.
7. **Benefits and measurable outcomes**

    Besides enhanced automation and minimal human intervention, some of the other benefits of the successful implementation of an agentic AI tool are:

    - Improved financial process efficiency and accuracy.
    - Minimal risk of compliance breaches with a proactive audit trail and process in place.
    - Lower operating costs and number of fraud cases make it easier to work with vendors and adapt to new requirements.
    - Automation of data extraction, validation and routing eliminates the need for manual work.
    - A faster and more reliable invoice journey with fewer challenges enables liquidity forecasting for the decision makers.
    - Quick activation of new suppliers with AI-based validation templates which supports special invoice formatting and contractual variations.

Select measurable outcomes of adopting agentic AI solutions for the client were:

Reduced manual processing time
85% decrease in manual processing time per invoice; from an initial spent of 2–3 days down to less than 10 minutes on an average.

Faster deployment
Time from pilot to production deployment was shortened to 1.5 weeks versus the standard 4–6 weeks for traditional automation rollouts.

Lower operational cost
Automation of data extraction, validation and routing processes reduced accounts payable operations costs by up to 60%.

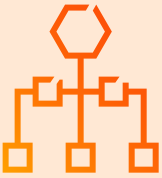

Minimal errors
A massive 85–90% dip in keyed errors meant reduced exposure to compliance penalties and payment disputes.

Standardised templates
The move to 5% standard templates from error-prone macros eliminated the need for extensive maintenance, reducing compliance incidents by 70%.

Centralised invoice repository
The centralised invoice repository was able to reduce compliance issues by 70%, providing a single source of truth regarding the vendors.

Enhanced vendor processes
The organisation was able to simplify onboarding processes and was also able to streamline invoice validation and processing.

Enhanced compliance and audit readiness
The solution helped automate compliance requirements with internal policies to ensure that contract terms and regulatory alignment are fulfilled.

Rapid deployment
The solution was quick to deploy and could go from pilot to a minimum viable product (MVP) in less than 30 days.

Looking ahead

Agentic AI is a game-changing advancement for organisations that can take invoice validation processes from passive legacy automations to a dynamic, intelligent orchestration. Adopting agentic AI offers tangible business benefits including

- quicker resolution of financial tasks
- improved vendor relationship
- higher audit preparedness
- better third-party risk management.

By shifting from rules-based automation to AI-driven orchestration, finance teams can focus on high-value, judgement tasks and rely on AI to process routine validation. Future use cases of agentic AI could include predictive vendor profiling, dynamic approval processes and conversational analytics; however, the focus of the C-suite and finance function should be on developing collaborative human-AI teams where AI is not a mere tool but an adaptive co-pilot that helps create smarter, faster and more resilient financial ecosystems.

Contributing to this article were **Anirudh Singh Rana**, **Sachin Parashar** and **Rodney D'Souza.**

About PwC

We help you build trust so you can boldly reinvent

At PwC, we help clients build trust and reinvent so they can turn complexity into competitive advantage. We're a tech-forward, people-empowered network with more than 370,000 people in 149 countries. Across assurance, tax and legal, deals and consulting we help build, accelerate and sustain momentum. Find out more at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

Contact us

Rajnil Mallik
Partner and AI GTM Leader
rajnil.mallik@pwc.com

Sumit Srivastav
Partner and Leader – Agentic Automation
sumit.srivastav@pwc.com

Dr. Indranil Mitra
Partner – iDAC (intelligent data, agents and cloud)
indranil.mitra@pwc.com

Ankit Garg
Partner – Risk and Fraud Analytics
garg.a.ankit@pwc.com

Contributors to this edition

Vishnupriya Sengupta
Hariprasad Gajapathy
Anirudh Singh Rana
Disha Thadani
Ruchika Uniyal
Sachin Parashar
Rodney D'Souza
Megha Adhikari

Editorial

Dion D'Souza and **Rubina Malhotra**

Design

Kirtika Saxena and **Harshpal Singh**

pwc.in

HS/August 2025 - M&C 47822