

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.

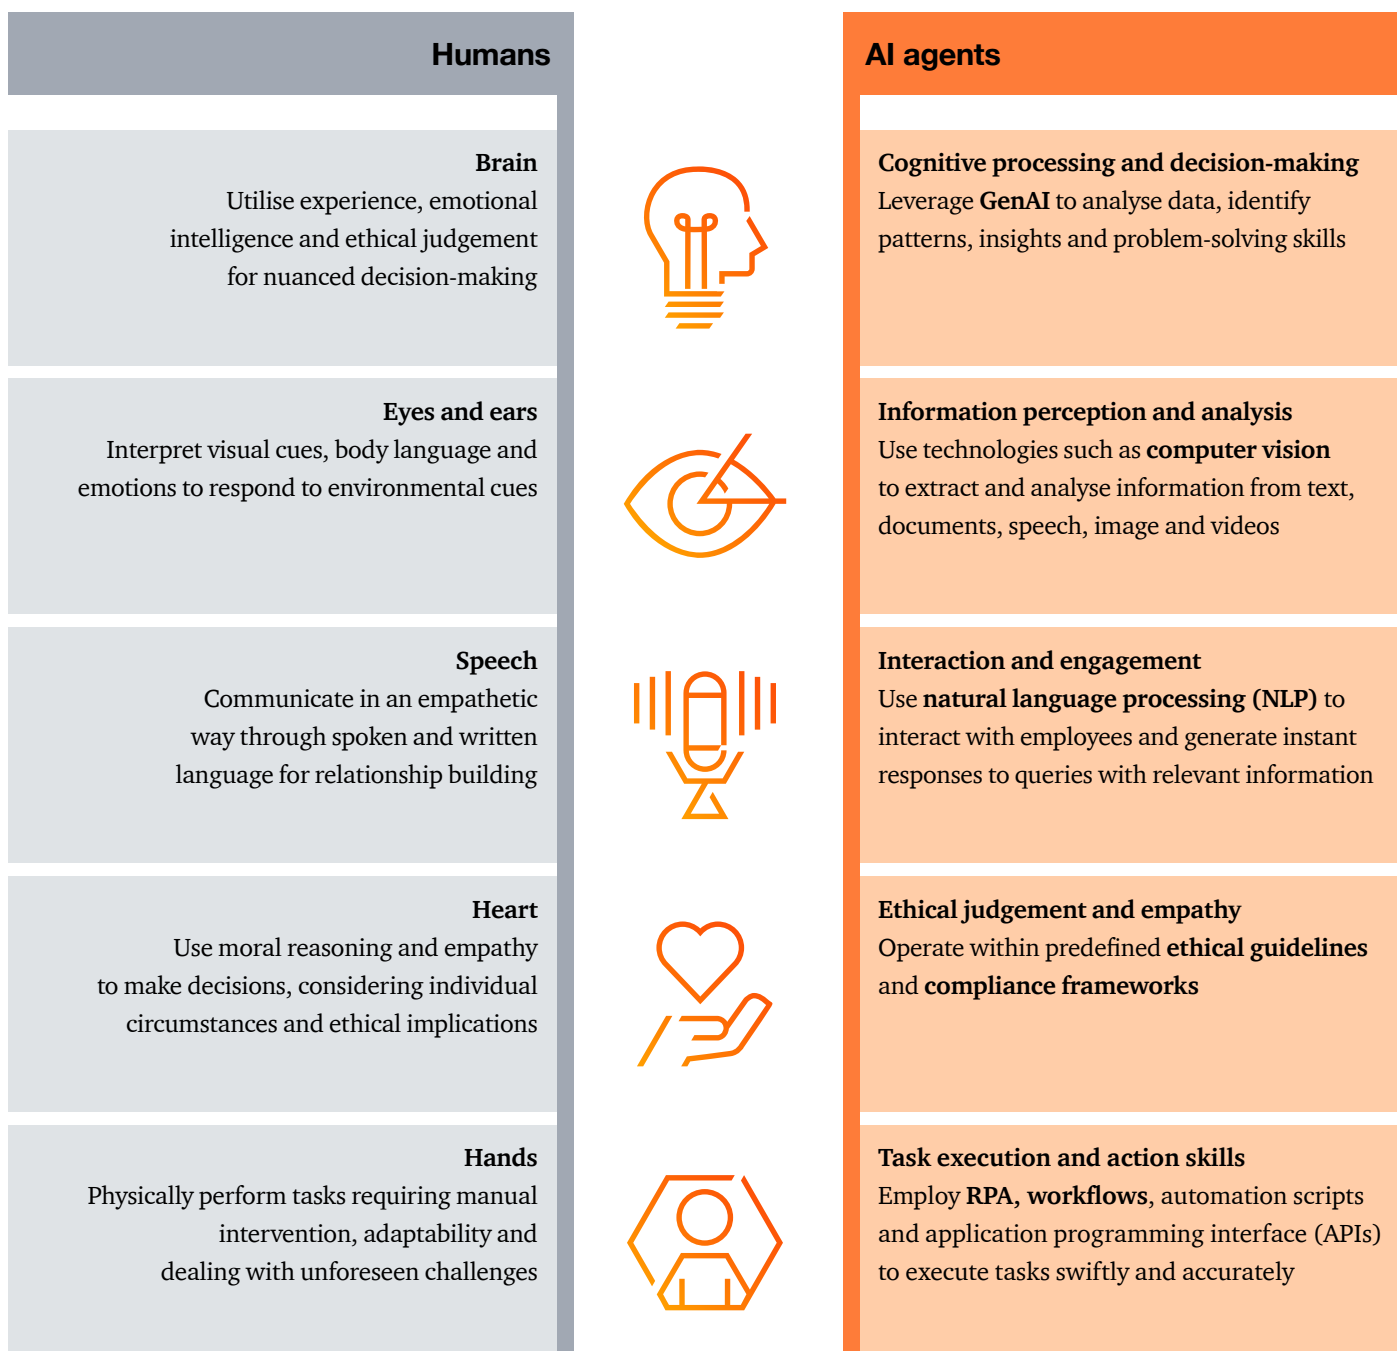


Combining human creativity and empathy with AI's efficiency can result in a transformative force which can reshape leadership, build trust and drive adaptive change across enterprises.

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



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.



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.

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.

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 Cycle¹ 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



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. | <ul style="list-style-type: none">Map high-impact use cases to business KPIs.Secure C-suite sponsorship and budget. |
| Govern | Establish robust, responsible AI guardrails. | <ul style="list-style-type: none">Define ethics, security and compliance policies.Implement risk assessments and audits. |
| Enable | Build technical and data foundations. | <ul style="list-style-type: none">Deploy a modular, multi-agent orchestration platform.Create AI-ready data pipelines. |
| Nurture | Cultivate talent and culture for human–agent collaboration. | <ul style="list-style-type: none">Launch targeted upskilling and change management programmes.Communicate agent roles clearly. |
| Transform | Pilot, measure, and continuously refine agent workflows. | <ul style="list-style-type: none">Run rapid sprints in controlled environments.Track accuracy, ROI and user satisfaction. |
| Scale | Expand proven agents across the enterprise responsibly. | <ul style="list-style-type: none">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

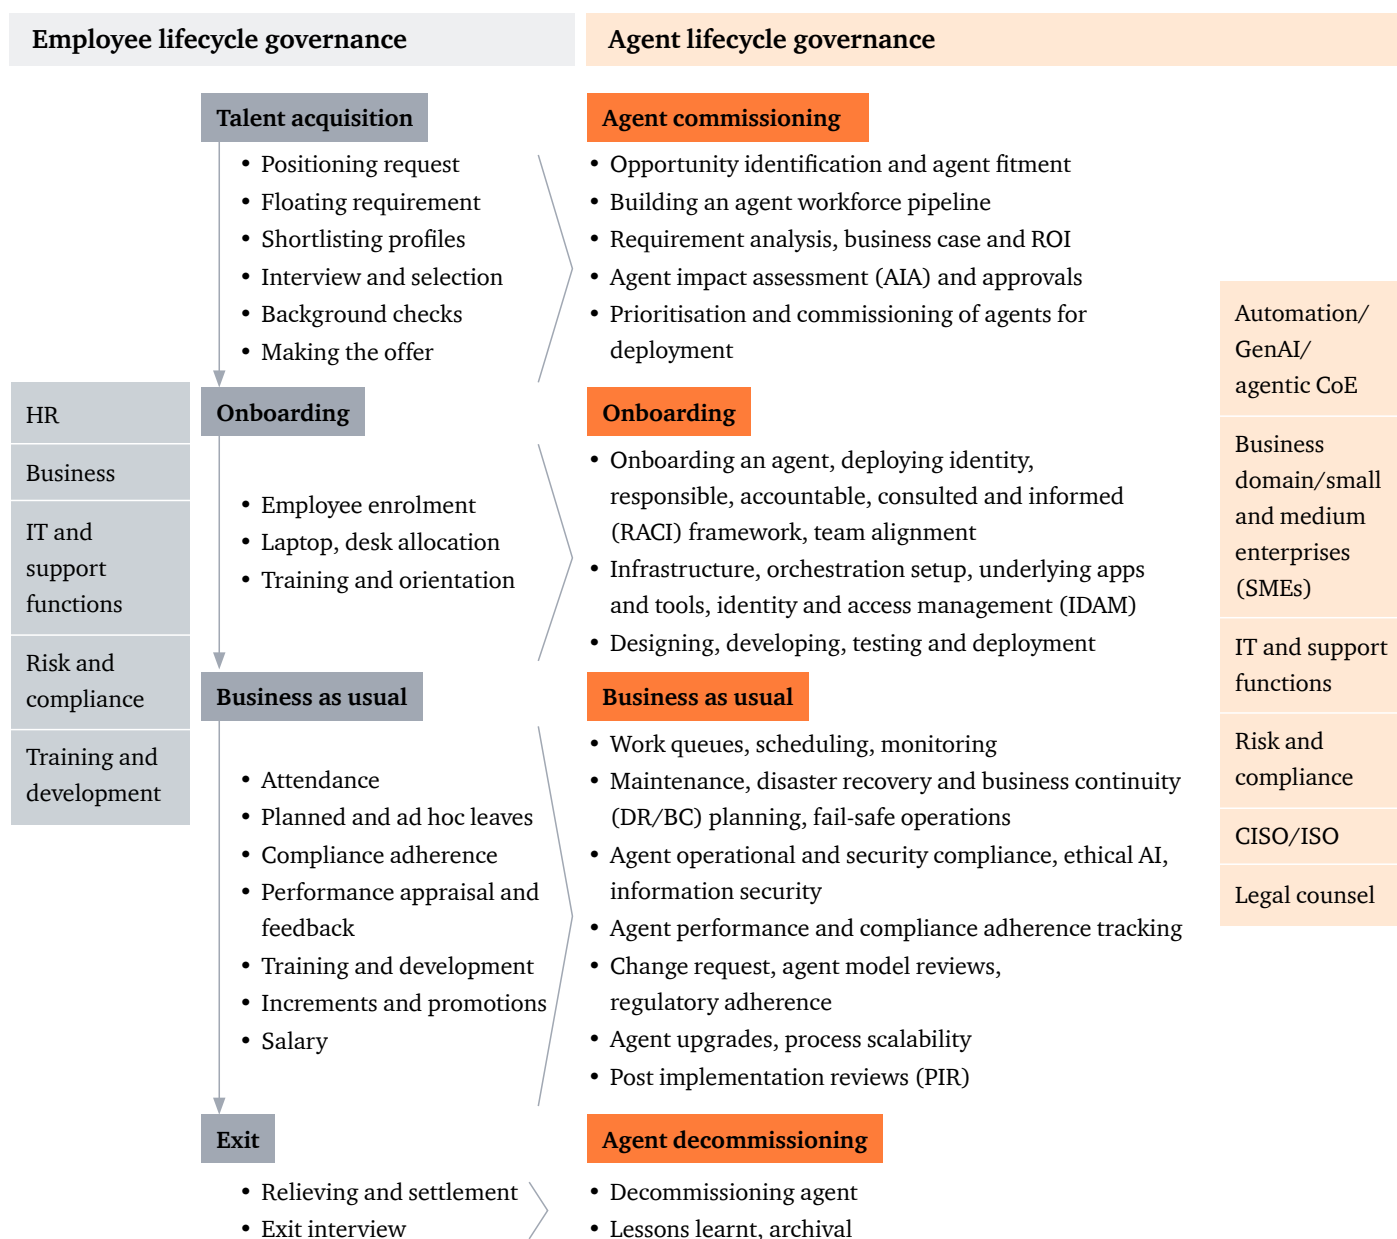
While GenAI tools have traditionally created content, generated predictions and delivered insights based on human prompts, agentic AI introduces a new technological frontier – systems which can think and act independently. This elevated autonomy not only transforms AI agents into proactive collaborators 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

| | |
|---|---|
|  | <p>Process governance</p> <p>Key tollgates in the agentic lifecycle from envisioning to monitoring of agents</p> |
|  | <p>Data privacy and security</p> <p>Protection of individual privacy and data while adapting to operational and environmental changes</p> |
|  | <p>Monitoring and reporting</p> <p>Key metrics to be tracked and reported for stakeholder and regulatory compliance</p> |
|  | <p>Human autonomy</p> <p>Upskilling required to understand, implement and monitor agents</p> |
|  | <p>Accountability and compliance</p> <p>Accountability for decisions made by agents to ensure compliance with legal and ethical guidelines</p> |
|  | <p>Fairness and bias</p> <p>Checking and improving the AI model's performance to ensure the decisions are fair and just</p> |
|  | <p>Transparency</p> <p>Agent design and decision-making process to be accessible to all relevant stakeholders</p> |

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.