pwc

One agent to rule them all

We build the agentic foundation, so you can unlock new potential.

PwC Switzerland in collaboration with IMD

IMD / Real learning Real impact

Table of contents

"This is not just another whitepaper to explain how cool AI is; rather, it provides guidance for leaders on how to build upon solid data foundations to enable effective implementation."

Matthias Leybold
Partner Technology & Data, PwC Switzerland

01

Executive summary

A new wave of opportunity is emerging for businesses ready to embrace change – to outthink challenges, outpace disruption, and achieve lasting impact. Artificial intelligence (AI) is entering a new chapter. What began with generative models producing text and images is now evolving into autonomous agents – intelligent systems that can perceive, decide, act, and learn with minimal human intervention. These agents are already reshaping how organisations operate, encouraging business leaders to reimagine not only their technology stack, but also their organisational structures, governance, and decision-making processes.

In earlier stages of AI maturity, the focus was on building a strong data foundation – a priority that remains essential in today's agent-driven environment. But expectations have grown. AI agents introduce new layers of complexity and opportunity, calling for not just technical readiness, but also strategic alignment, cross-functional coordination, and responsible oversight.

This whitepaper explores this evolution and how organisations can adapt to – and benefit from – a new AI era, where AI agents are becoming an integral part of business operations and everyday life, shaping new expectations from customers and employees alike. We explore these shifts so you can move forward with clarity and confidence – and outperform in a world where adaptability defines success.

Data remains the essential enabler. In the age of agentic AI, where intelligent systems act autonomously, success increasingly depends on high-quality, well-governed data. This paper builds on the evolution from One data strategy to rule them all – PwC's previous whitepaper on enterprise data foundations – to empowering your enterprise with agentic AI – a necessary upgrade for a rapidly changing world. The focus moves from managing static data assets to orchestrating dynamic, intelligent agents across the enterprise.

To examine this transformation from both a strategic and operational angle, PwC Switzerland partnered with IMD's leading AI professors Amit Joshi and José Parra Moyano. While PwC brings hands-on insights from real-world AI deployments, IMD's experts add a perspective grounded in ongoing dialogue with C-level executives navigating digital transformation.

Matthias Leybold, Partner Technology & Data, PwC Switzerland

Joscha Milinski, Partner Technology & Data, PwC Switzerland

Amit Joshi, Professor of AI Analytics and Marketing Strategy, IMD

José Parra Moyano, Professor of Digital Strategy, IMD

Together, we offer a balanced view of what AI agents are, how they differ from traditional AI tools, and what it takes to deploy them responsibly and effectively. The paper outlines key concepts, practical use cases, and actionable recommendations – not just to prepare for what's coming, but to shape it and realise real business outcomes.

This isn't a story about replacing people. It's about enabling them – to make better decisions, to be more creative, and to lead in ways that are fit for the future. And by working together, we can turn today's potential into tomorrow's progress, helping you stay ahead and lead with purpose.

02

The evolution of AI: from early promise to enterprise reality

To understand the current impact of AI, it helps to recall how far we have come. The quest for artificial intelligence began with symbolic AI in the 1950s and 1960s, when pioneers such as Alan Turing envisioned thinking machines and developed early programs that could follow logical rules. Initial expectations were high – too high, as reality proved. The field endured "AI winters" in the 1970s and late 1980s, when grand promises fell short and funding dried up. Nonetheless, crucial groundwork was laid during these periods: the emergence of expert systems in the 1980s and the development of machine learning approaches in the 1990s quietly kept AI progress alive. By the 2000s, steady advances in algorithms and computing power marked the beginning of a renewed era. Milestones such as IBM's Deep Blue and Watson, followed by the deep learning boom after 2010, demonstrated the potential of artificial neural networks – particularly as they began to outperform humans in tasks such as image and speech recognition.

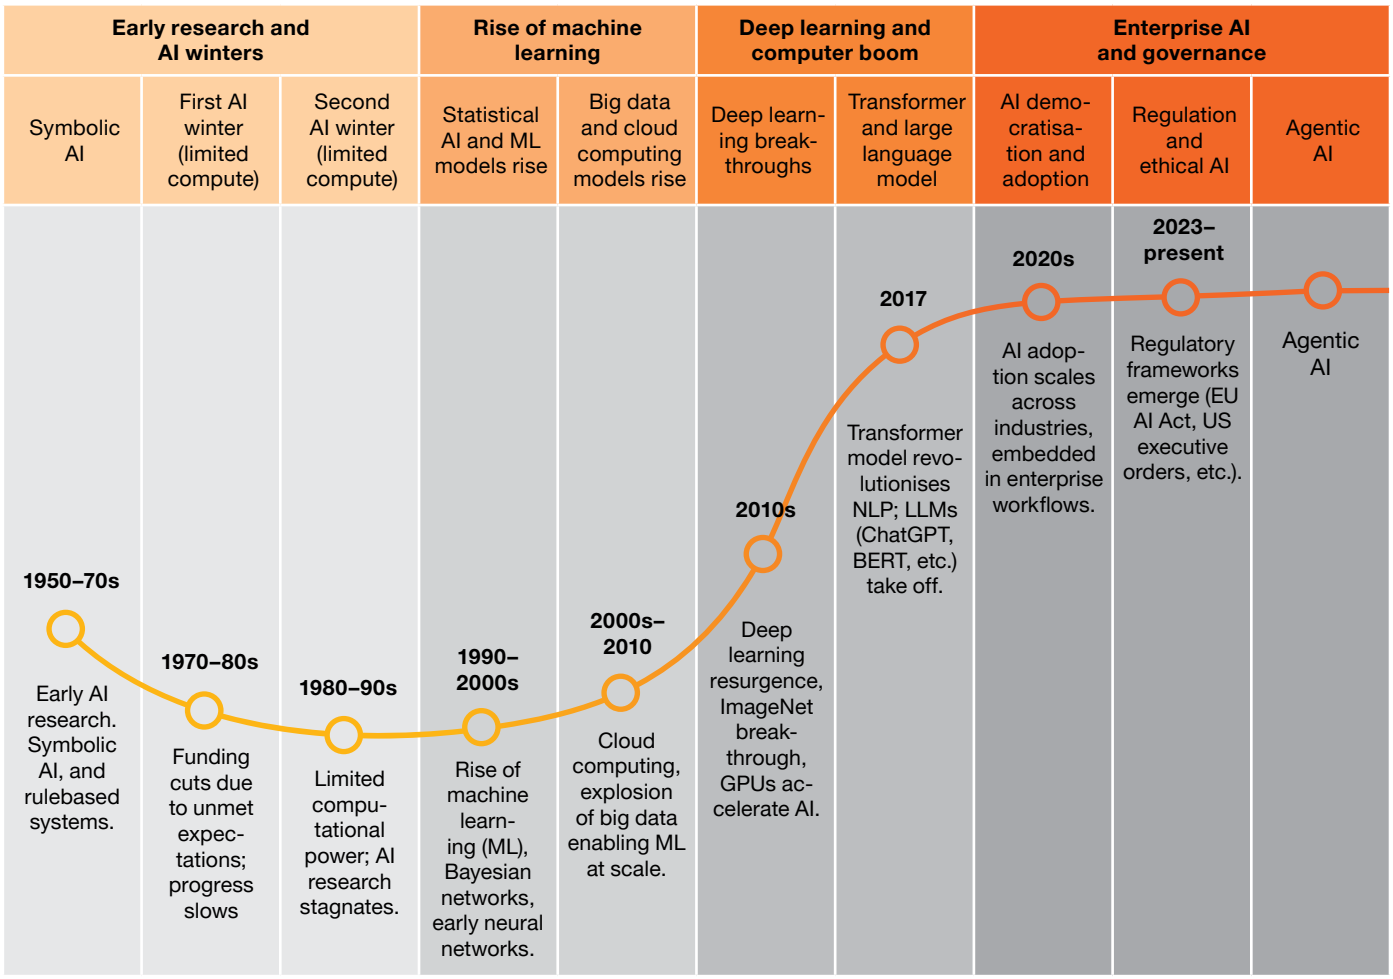

"GenAI is simply the world's best autocomplete."

Amit Joshi
Professor of AI Analytics and Marketing Strategy, IMD

A pivotal breakthrough came in 2017 with the introduction of the Transformer model, presented in the paper Attention Is All You Need, authored by eight researchers at Google. This architecture reinvented how AI systems process language, using self-attention mechanisms to handle words in parallel rather than sequentially. The impact was immediate and far-reaching: Google's BERT (Bidirectional Encoder Representations from Transformers), introduced in 2018, could understand language context with unprecedented accuracy, while OpenAI's GPT series scaled up to billions of parameters, culminating in GPT-3 in 2020. Each iteration demonstrated extraordinary leaps in the ability of AI to generate human-like text and identify patterns in vast datasets. More recent reasoning models such as O1 and O3 have further enhanced AI's ability to simulate human-like reasoning and decision-making. Building on these advances, AI has captured the attention of both the public and businesses – revealing its potential to converse, create, and assist at a level that increasingly feels like practical reality rather than science fiction.

Today, AI stands on the cusp of full-scale enterprise transformation, moving beyond isolated innovation projects. The narrative has shifted from viewing AI as a curiosity to recognising it as a strategic, integrated capability within leading organisations. Governments and regulators are responding by introducing guidelines to ensure that the growth of AI is accompanied by accountability and oversight. These guardrails signal that AI has matured from laboratory experiments into an enterprise reality – one that must be managed with the same rigour as any other mission-critical asset.

"AI agents are robotic process automation for unstructured data."

Amit Joshi
Professor of AI Analytics and Marketing Strategy, IMD

As organisations grow comfortable with AI performing discrete tasks, a new paradigm is emerging: AI agents capable of autonomously executing complex sequences of actions. Unlike a single machine learning model that might, for example, predict monthly sales, an AI agent functions more like a virtual team member – it perceives its environment, makes decisions, takes actions, and learns from the outcomes. AI agents operate at the intersection of AI, automation, and system integration, offering organisations an unprecedented opportunity to expand their AI capabilities and transform how tasks are performed and complex challenges are addressed across the enterprise.

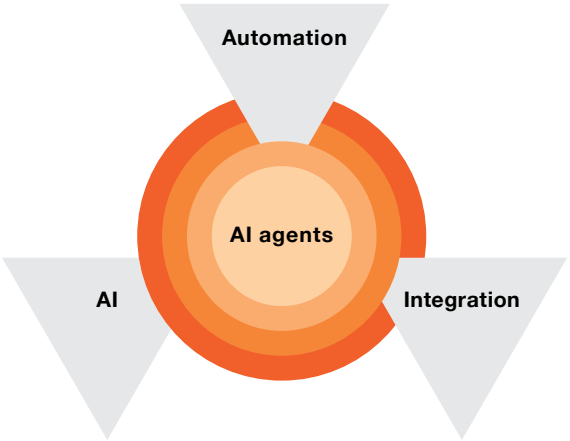

Given the transformative impact of AI agents, this whitepaper focuses exclusively on their role –providing organisations with the keys to unlock their potential, starting with how to strategically shape a data strategy to enable their adoption.

03

A data strategy framework that still rules them all

The foundation for successful data strategies has not changed; AI agents are merely the visible part of the iceberg. While PwC's enterprise data strategy (EDS) framework remains valid, AI agents have expanded the role of AI within organisations – and in doing so, have added complexity to the strategic decision-making process. Understanding how AI agents should be strategically targeted is key to realising their full potential.

To learn more about PwC's EDS, please refer to our previous whitepaper A comparative perspective on data strategies.

To address the complexity and variety of data, business processes, and data governance requirements, PwC's EDS framework consists of five main dimensions and 27 sub-dimensions. The five core dimensions of the data management framework are:

- Data governance
- Data processes and tools
- Data architecture and quality
- Data security and privacy
- Culture and people

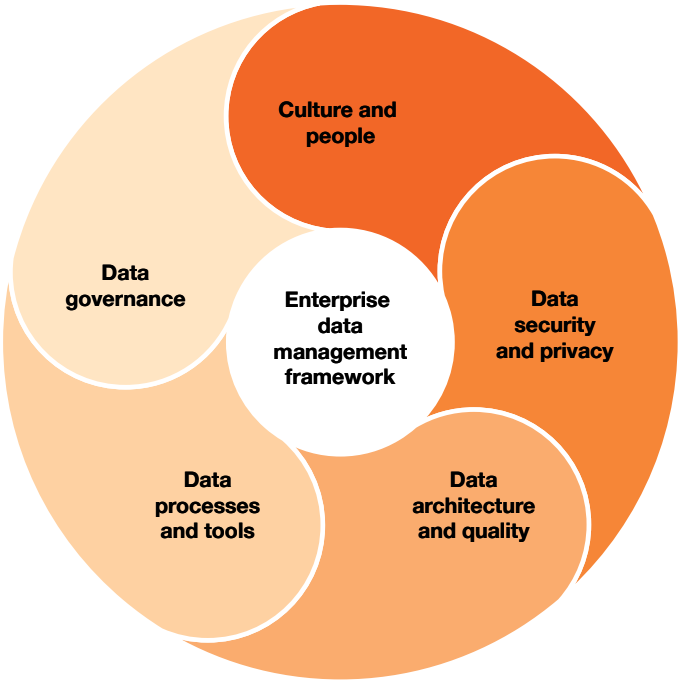

Before the emergence of AI agents, the EDS framework identified critical dimensions based on the maturity of the organisation and the nature of the targeted use case. As AI agents have broadened the scope of AI within organisations, these dimensions have become more interdependent. One could argue that "all EDS dimensions are equally important" as agents have the potential to impact the entire organisation. While this holds true in most situations, companies must still prioritise to ensure efficient allocation of resources and effort.

To account for this, it is worth considering how AI agents can be promoted across an organisation. In our view, AI agents can be driven from either an enterprise or a citizen-led perspective – or both. The enterprise perspective follows a "top-down" approach, where specific use cases are prioritised. In contrast, the citizen-led perspective adopts a "bottom-up" approach, focusing on empowering individual users. Here, federation and broad democratisation – or, in other words, widespread adoption – are the driving factors. These two approaches are not mutually exclusive and should ideally coexist. However, depending on the chosen perspective, organisations can shape their data strategy with a targeted focus on how to best capitalise on the potential of AI agents.

Enterprise driven
Enterprise driven AI agents focus on E2E use cases requiring dedicated teams of specialists often located centrally:

- **Use case based identification**
- **Single AI agent holds big value unlocks**
- **Longer lead time to implementation**
- **Each AI agent serves multiple users**
- **Enterprise governed**

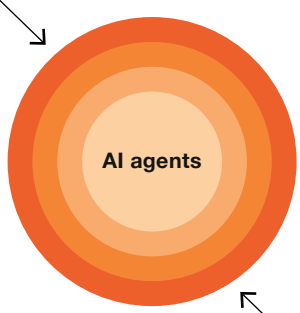

Citizen driven
Local, highly specialised AI agents that cater to a single or a specified set of users. Commonly built by the users themselves:

- **Situational and ad-hoc identification**
- **Single AI agent holds smaller value unlocks**
- **Short lead time to implementation**
- **Each agent serves single/limited users**
- **Locally governed**

Case studies

To give examples of the two approaches described above, we describe two case studies based on the citizen-driven and enterprise-driven approaches. These case studies demonstrate how AI agents can be effectively implemented in different organizational approaches.

Citizen-driven approach

Benefits:

The Finance Commenting Agent streamlines the reporting process, making it more efficient and comprehensive.

In the financial industry, consolidating the quarterly reports of companies within a portfolio to obtain a yearly overview consumes time from financial analysts. Time that could be better spent on value-adding tasks such as analysing the portfolio's performance and shaping the best advice for investors. The Finance Commenting Agent is a strong example of the citizen-driven approach. The idea was brought up by the analyst and was quickly built, tested and implemented in CoPilot Studio. Its main goal is to streamline the creation of consolidated quarterly reports, seamlessly integrating with SharePoint for efficient document management.

When activated, the agent guides the analyst through selecting the relevant year(s), quarter(s), and financial position(s) for the report. It then extracts both structured and unstructured data from subcompany reports (including comments), generates a consolidated financial report, and offers multiple export options – ultimately saving valuable time by automating a previously tedious task.

Enterprise-driven approach

Benefits:

The Vendor Master AI Agent increases accuracy and efficiency in vendor management, minimises manual errors, and supports better decision-making through timely and reliable data.

Some topics – such as master data management – are too broad and complex to be tackled by citizens alone and instead require a coordinated, enterprise-driven approach. The Vendor Master AI Agent is a good illustration of this. Creating vendor master data is often associated with rigid, hardcoded automation and manual copy-paste tasks from documents such as quotes or invoices. At best, once logged, this data is buried deep within a master data tool; at worst, it is stored in a spreadsheet on a shared drive.

Acting as a companion for managing vendor data, the Vendor Master AI Agent solves both problems. It automates data capture through easy document uploads via chat, flags discrepancies for reviews (e.g. potential vendor duplicates), and enables users to find vendors using both structured and unstructured descriptions (e.g. "retailer from Delaware"). Lastly, the agent continuously improves by incorporating user feedback and is scalable for integration with various business processes, providing efficient and adaptive support for vendor management.

04

Key factors for an AI agent-ruled business world

Today's leaders tend to focus on an internal perspective, often overlooking critical outside influences that affect their AI journey. With the rise of AI agents, a set of external forces is placing increasing pressure on the enterprise data strategy. Regulation, cybersecurity, emerging technology trends, and ethical considerations are among the key external factors we have identified that will have an impact on the subsequent data strategy.

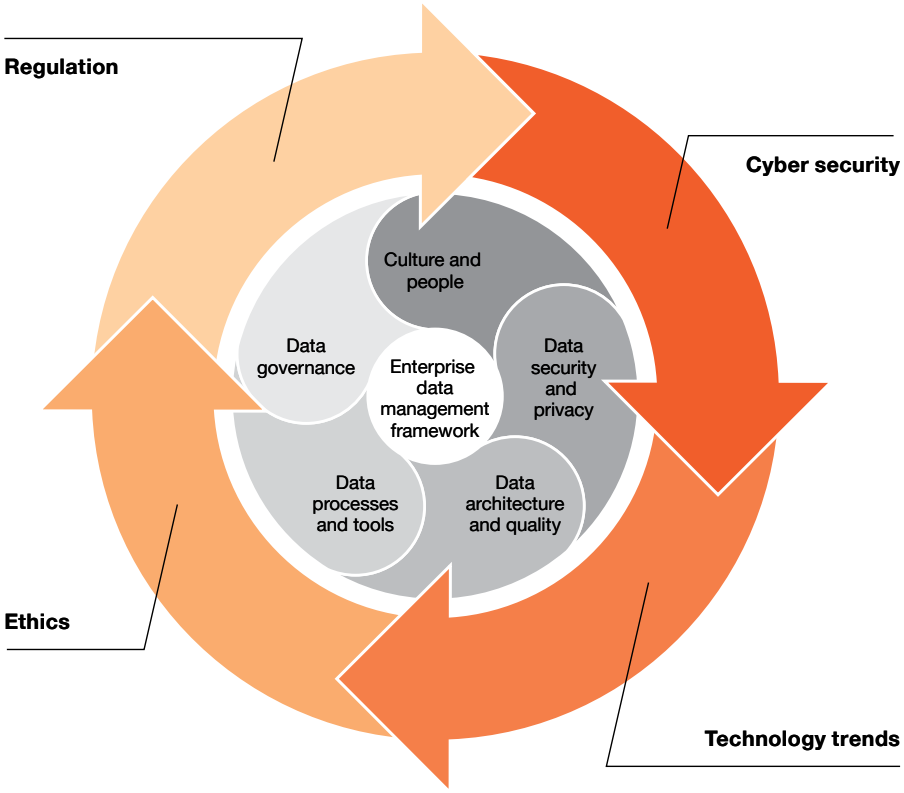

Regulation

While few policies target AI agents directly, they fall under broader regulations on data protection, algorithmic transparency, and digital accountability. Global approaches vary: the EU prioritises strict governance with frameworks such as the AI Act, GDPR, and Data Act; the US and China favour innovation through looser, market-driven controls; other regions adopt more adaptive, sector-specific strategies. In this dynamic landscape, regulatory agility is crucial. Business leaders must track emerging frameworks, anticipate compliance requirements, and align their AI strategies accordingly.

"The winners won't prevail because of better technology. They'll likely succeed by rethinking the nature of work – and what that means for the organisation."

PwC US, "AI agents can reimagine the future of work, your workforce and workers", 2025

Cybersecurity and privacy

AI agents intensify cybersecurity challenges by acting independently and accessing critical data and systems. This expands the potential attack surface and increases exposure to new risks such as prompt injection, model inversion, or data poisoning. As AI agents gain access to sensitive data and perform actions on behalf of users, security and privacy controls must evolve accordingly. At the same time, agents offer opportunities to improve cyber defence, such as automated threat detection and response. The effective use of AI agents requires robust safeguards: clearly defined system privileges, data protection mechanisms, and human oversight to ensure safe and reliable performance.

Technology trends

AI agents are shaped by rapid developments in both software and business models. One key trend is the growing momentum behind open-source AI – making powerful models more accessible, customisable, and affordable. While this fosters faster innovation, it also introduces challenges in terms of integration, security, and governance. At the same time, changing licensing structures for proprietary tools can drive up costs and create lock-in risks. To stay competitive, organisations need to manage these dynamics strategically, balancing flexibility with governance and long-term sustainability.

Ethics

Autonomous agents amplify long-standing ethical concerns. Biased training data and hallucinated outputs can cause harm at scale if left unchecked. Mitigating these risks requires ongoing model improvement and careful validation to ensure that outputs remain accurate and trustworthy. Their growing autonomy also raises questions about workforce impact and accountability. Leaders must actively manage fairness, transparency, and inclusion, while ensuring humans remain in control of strategic decisions. Clear communication and investment in AI literacy will be essential for building trust – internally and externally.

05

From experiment to impact: turning AI pilots into measurable returns

We are experiencing an unprecedented rate of innovation, where the "tomorrow" of today and the next iteration is just around the corner. Organisations now need to manage technology investments in an increasingly dynamic environment. This poses a critical challenge to leaders: should they wait and buy when off-the-shelf solutions come to market, or should they invest and build the solutions required today? Well-versed and experienced leaders understand that this is not a trade-off or a binary decision – it is a balance to be managed.

Standing on the sidelines equals risk. Companies today are aware and often concerned about being outpaced, missing out on comparative advantages, and ultimately lacking the ability to capitalise on the AI promise. At the same time, staying at the forefront also carries risk, as significant investments may lose value with each new evolution. This places emphasis on ROI management: leaders should not be afraid to invest in AI agents but instead focus on ensuring those investments deliver returns quickly.

Yielding a positive ROI may sound obvious, but it is not enough. Achieving a quick return is arguably even more important. Leaders who focus on this will not only reap the benefits of AI, they will also be well-positioned when the next evolution arrives. As they enhance their maturity through the development of their own agents, they are able to test, fail fast, iterate, and succeed. With that enhanced maturity, when technology providers roll out off-the-shelf solutions, these organisations are best positioned to assess and configure swiftly to replicate functionality from their previous agents (e.g. selected LLM models, parameter settings, data sources integration, etc.). Lastly, since their in-house agents have already delivered value, there will be no heartbreak in the decision to decommission them and move on to the next generation of agents.

06

Enterprise functions drive innovation

The previous sections presented two approaches to driving AI agent adoption (enterprise and citizen), along with the key influencing factors and the strategic balance between them. This may create the impression that benefiting from AI agents is a daunting task, requiring extensive upfront analysis, strategy, and planning. However, our research and industry experience show that the use and adoption of AI agents – and the tools to build them – are already spreading and accelerating. Therefore, organisations must carefully balance the adoption of AI agents with a strategic approach to steering the technology. Those who do, will be able to continuously evolve and capitalise on the promise of AI. Those who do not, risk building an unstable and uncontrollable asset, missing out on its full potential.

In most companies, functions have the greatest impact on shaping the enterprise data strategy. Each function has distinct leadership, clear scope and integration points, and manages its own budget. As such, they are best positioned to improve the maturity of the strategy and to define how AI agents should be deployed and used effectively.

For simplicity and ease of application, we have grouped functions into the following categories. Depending on the company context, some may be split, combined, or new ones could be added:

- **Strategy**
- **HR**
- **Marketing and sales**
- **Finance and accounting**
- **Service and product development provisioning**
- **IT**
- **Legal and compliance**

While functions control the approach and make “build or buy” decisions, they are not equally affected by the external factors introduced above. Some manage proprietary knowledge and processes that are instrumental to the company’s differentiation, while others operate in heavily regulated environments or handle highly sensitive data.

We have discussed the impact of external factors on each function with many of our clients, debated them in our classrooms and speaker panels, and arrived at the following conclusion (see next page Source: PwC internal research):

	Strategy	HR	Marketing and Sales	Finance and accounting	Service and product	IT	Legal and compliance	**Total**
Regulation	1	4	2	3	3	2	4	19
Ethics	1	4	3	1	3	1	2	15
Technology trends	3	3	3	3	3	3	3	21
Cyber security	4	4	4	4	4	4	4	28
Total	9	15	12	11	13	10	13	

Legend: 1 = very low impact
4 = biggest impact

Before diving into the differentiating factors, it is important to highlight that cybersecurity received the highest score across all functions. Given the ability of AI agents to make autonomous decisions, ensuring security and protection against malicious intent is paramount. This may seem obvious for externally facing AI agents such as a social media agent tasked with amplifying employees' posts on LinkedIn. Imagine an agent commenting on a new joiner's post with: "Welcome onboard, looking forward to working together – and thanks for helping us meet our diversity and inclusion quotas". But sound cybersecurity it is equally important for internal service desk AI agents that help employees reset their passwords. Given our overall reliance on laptops and phones, hackers could easily bring a company's operations to a grinding halt by taking over such an agent.

Similarly, we observe that the technology trend has a consistent impact across all functions. While slightly less critical than cybersecurity, the pace of innovation requires continuous monitoring and adaptation. These two factors impact functions throughout the organisation and should be considered of comparable importance.

Regulation and ethics are therefore the key differentiating factors. AI does not operate in a lawless vacuum – data privacy laws such as GDPR grant individuals rights over their data, and regulators worldwide are scrutinizing how AI algorithms make decisions that affect consumers. The EU AI Act, for instance, not only classifies high-risk AI systems but also explicitly bans certain AI use cases deemed too dangerous, such as social scoring of citizens or real-time biometric identification in public without proper oversight. It also requires that AI-driven decisions be transparent and explainable to users, echoing the principle that individuals have a right to understand how an agent's underlying algorithm affects them.

Intuitively, the strategy and IT functions are the least impacted, allowing AI agents to operate with more freedom and requiring less oversight for deployment. In contrast, the HR function is the most affected due to data protection regulations (e.g. GDPR, EU Data Act) and the significant implications for the workforce. For example, an AI agent used for candidate matching must be free from bias, and the sensitivity of personal data places strict limitations on the scope of such an agent.

Case study

AI agents for candidate matching

In a recruitment industry marked by increased competition for talent and access to a much wider pool of candidates through remote and hybrid work, AI agents are used to match candidate profiles with job descriptions, optimising the recruitment process by ensuring the best fit between candidates and roles. This approach helps recruiters and hiring managers efficiently reduce the manual effort and time between a job posting and the selection of new hires. Candidates benefit from faster feedback on their applications and the reassurance that an agent is helping to reduce bias.

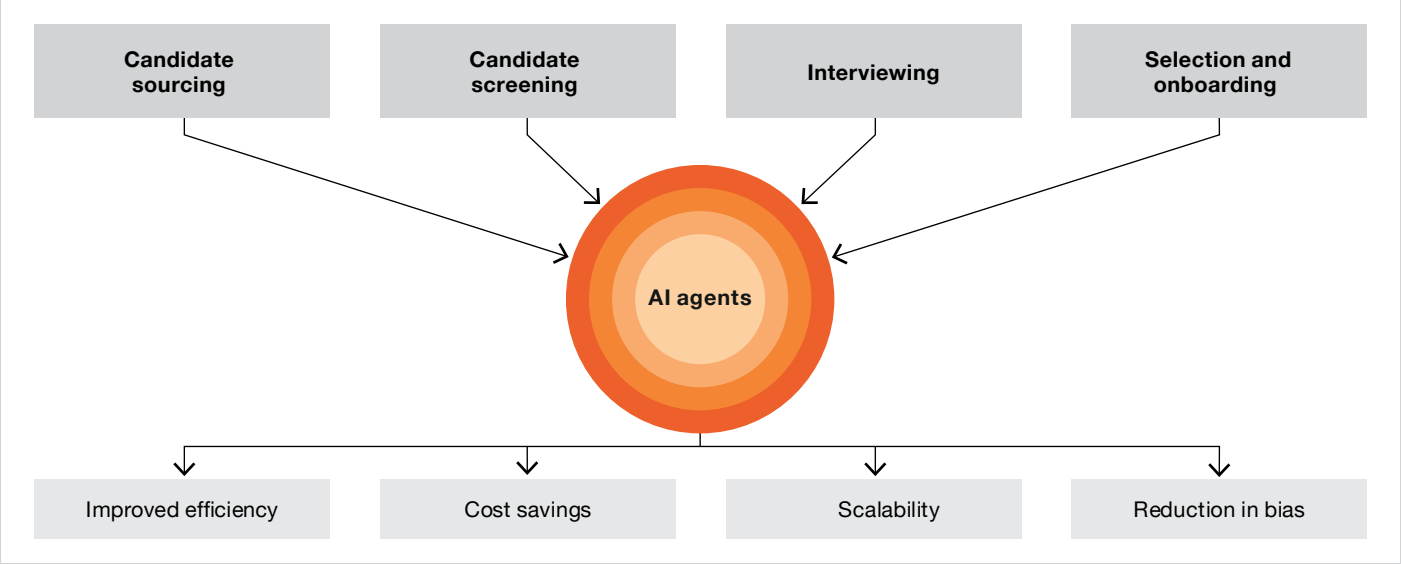

- **Candidate sourcing:** AI agents can automate the sourcing process by scanning job boards and social media for suitable candidates, reaching out to passive talent, and managing candidate databases. They can also analyse past sourcing data to refine strategies and expand reach.
- **Indidate screening:** AI agents can expedite screening by using algorithms to evaluate resumes and applications for key skills, experience, and qualifications. Natural language processing helps identify relevant skills and rank candidates accordingly, making the screening process faster and more efficient.
- **Interviewing:** They assist in scheduling interviews by finding ideal times for both interviewer and candidate, automating reminders, and preparing tailored interview questions based on candidate profiles. They can also support preliminary interviews via chatbots, providing initial assessments.
- **Selection and onboarding:** AI agents improve decision-making by analysing data collected throughout the recruitment process to highlight the best candidates. For onboarding, they can manage workflows, ensure completion of necessary documentation, and personalise the onboarding experience based on the candidate's needs and role-specific requirements.

Attention points

While the benefits are clear for both candidates and recruiters, this is a good example of how regulation and ethics become key differentiating factors:

- Multiple **regulations** come into play – from employment laws to data privacy and AI-specific regulations. These require clear protocols for data usage and storage, as well as transparency with candidates about how their data is processed and used by AI systems.
- While far from perfect, HR teams have been working for years to improve the **ethical standards** of their recruitment process (e.g. removing pictures and names from CVs, providing bias training for anyone involved in the recruitment process). Given their inherent bias, agents should be regularly tested and audited to allow for corrective measures and prevent discrimination, thus fostering the intended diversity and inclusion benefits.

07

The impact of external factors shapes the strategy

Through our conversations, we have also identified a correlation between the impact score and the favoured approach to AI agent adoption: the higher the score, the more enterprise driven use cases are preferred to deploy AI agents. This preference is justified by the higher need for control, governance, cost, and time to deploy, factors that are counter-balanced by the higher benefit potential of a single AI agent. Enterprise-driven AI agents require dedicated teams, including specialists, who focus on specific use cases that deliver end-to-end benefits across teams, functions and the organisations as a whole.

The citizen-driven approach aims to deliver many small wins. It is a systematic method of generating value through an increasing number of more engaging experiences, higher-revenue products and services, more efficient workforces, and more productive workflows. More specifically, AI agents in the citizen-driven approach are tailored solutions – often built by the users themselves – acting as close companions that make everyday tasks more efficient and enable the users to refocus on value-generating activities. These AI agents operate within governed environments but with more freedom than those developed under the enterprise-driven perspective. They operate within teams, and in some cases serve only an individual user, requiring broad upskilling of the teams responsible for building AI agents.

As mentioned before, these approaches are not mutually exclusive and should not be considered in isolation: regardless of the score, functions can derive value and return on investment from both approaches, which can – and should – coexist. However, they may not carry the same priority. Depending on the score, the operating model and the initiatives launched should be shaped to fit the purpose and ensure that investments are appropriately targeted.

At the intersection of the two approaches, citizen-built agents that prove valuable across users or teams are singled out and further invested in to scale beyond their initial use. Conversely, widely used enterprise agents may be adapted or replicated to address more specific needs. To do this, organisations need to establish an operating model built around a marketplace, where agents from both approaches can be published, consumed, and further enhanced through feedback loops.

08

Looking beyond the known: how AI agents determine C-suite agendas

One of the most significant misunderstandings about AI agents is the idea that they will soon render humans obsolete. While agents can act autonomously and make decisions within defined parameters, the architecture and governance that guide their behaviour remain in human hands. The notion of a "human in the loop" is not only a regulatory requirement in many jurisdictions but also a strategic and practical necessity. While the technology is advancing at a rapid pace, it is important to remember that machines are not ruling us — people are still at the helm.

Human judgment is still irreplaceable when it comes to setting priorities, resolving ambiguities, interpreting nuanced contexts, and making ethical decisions. This is especially important when addressing questions that have no clear right or wrong answers, when limited or no data is available, or when uncertainty – rather than measurable risk – dominates the context. And this is precisely where business leaders operate.

In this context, business leaders, by the very nature of their role, will not delegate control – they will "simply" expand their toolkit. A well-functioning AI agent ecosystem will enhance the capabilities of leaders and does not replace them. CEOs, CFOs, and other top executives must steer this ecosystem with the same rigour they apply to finance, operations, or risk management.

What about the executers, not the executives?

Not everyone is a CEO, and thankfully so. Most people work by executing tasks rather than by designing them (let alone developing strategies or long-term plans). Regardless of the extent to which AI agents take over execution tasks, the big (and classic) question for organisations will be: **how do we grow, not merely by getting a bigger slice of the cake, but by baking a larger one?** Answering this question requires creativity, the ability to handle uncertainty, and to navigate a high degree of ambiguity. Therefore, it is a question best suited for AI-enhanced humans, not for AI agents alone.

"We show you how to leverage data, so you can drive AI innovation."

Joscha Milinski
Partner Technology & Data,
PwC Switzerland

Successful organisations will be the ones that convert **executers** into **explorers,** equipping them with the latest cognitive technology to tap into the unknown, ask better questions, and tackle the quest for growth in ways that would not be possible without AI enhancement.

While this view may be criticised for reflecting traits of voluntarism – the belief that **will,** and not **rationality,** is the primary driver of human behaviour and decision-making – it is important to remember that humans have consistently demonstrated a remarkable capacity for intuition, creativity, and adaptability This capacity is hard-wired into us all. Therefore, it is not far-fetched to suggest that organisations investing in AI-augmented capabilities, paired with strong change management and adoption efforts, will see their executers transform into explorers – professionals empowered to ask deeper questions, innovate boldly, and adapt strategies on the fly. This synergy between human insight and machine efficiency can create a dynamic environment where growth is not measured only by scale but by the ability to reinvent and expand the very boundaries of possibility.

Functions drive the agent agenda – but humans still rule them all

As described earlier in this whitepaper, enterprise functions are best positioned to have a meaningful impact on AI adoption and scale its impact. Each function owns its data, its decision-making processes, and increasingly, its agents. However, this distributed power does not eliminate the need for overarching coordination. The metaphor of the CEO (or perhaps a CAO, Chief Agent Officer) as the "agent of agents" is both symbolic and operationally relevant. Just as a functional AI agent is responsible for a specific domain – such as automating compliance processes or optimising media buys – this person must orchestrate the entire system. They set the strategic guardrails, ensure alignment, and arbitrate the trade-offs between competing interests.

Faced with this new agent architecture, CEOs and executive teams must rise to a new mandate. It is no longer sufficient to have a digital strategy alone; leaders must ensure that their broader AI strategy includes a clear view on the role of AI agents in driving growth. This involves making decisions about which agents to build versus buy, which functions to prioritise, what ethical guardrails to establish, how to monitor performance while ensuring cybersecurity and data protection, and – critically – how to upskill the human workforce to convert executors into explores who can collaborate effectively with AI and foster growth.

In this new environment, C-suite leaders are not just decision-makers. They become orchestrators of a hybrid workforce in which digital agents and human colleagues coexist, interact, and ideally, co-evolve. Just like any other high-performing team, the agent workforce requires guidance, governance, and clear accountability mechanisms.

A final word: Choosing to build, not to surrender

Many executives – and citizens – are understandably anxious about a future where machines appear to take over the workplace, decision-making, and even creative thinking. But this anxiety is nothing new. The fear of technology replacing people dates back at least two centuries. It was present during the Industrial Revolution, and resurfaced with the second, third, and fourth (and now the fifth) industrial revolutions.

"Let us go back to our desks – not to surrender, but to imagine and to build."

José Parra Moyano
Professor of Digital Strategy, IMD

What history teaches us is this: each revolution brings massive opportunities and humans are remarkably adaptive. Every technological disruption has eventually led to new forms of employment, new ways of creating value, and even new dreams. The emergence of AI agents is no different. Yes, they are powerful. Yes, they will change how we work and live. But they will not strip us of our agency. On the contrary, they offer us the chance to dream bigger, pursue moonshots, and redefine what we believe is achievable.

Let us not be lulled into complacency or paralysed by fear. Let us go back to our desks – not to surrender, but to imagine and to build – with ambition, with courage, and with responsibility. If AI agents are here to stay, let us ensure they serve a world that is more prosperous, more sustainable, and more inclusive – not despite our humanity, but because of it.

Bibliography

One data strategy to rule them all, PwC Switzerland, **https://www.pwc.ch/en/insights/digital/a-comparative-perspective-on-data-strategies.html**

Attention Is All You Need, Google, **https://research.google/pubs/attention-is-all-you-need/**

AI agents can reimagine the future of work, your workforce and workers, PwC US, **https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agents.html**

2025 AI Business Predictions, PwC US, **https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html**

Authors

Matthias Leybold
Partner Technology & Data,
PwC Switzerland
matthias.leybold@pwc.ch

Joscha Milinski
Partner Technology & Data,
PwC Switzerland
joscha.milinski@pwc.ch

Rodolphe Dehaspe
Manager Technology & Data,
PwC Switzerland
rodolphe.dehaspe@pwc.ch

Martin Gundersen
Manager Technology & Data,
PwC Switzerland
martin.g.gundersen@pwc.ch

Amit Joshi
Professor of AI, Analytics and
Marketing Strategy, IMD

José Parra Moyano, PhD
Professor of Digital Strategy, IMD

Find out more about our digital services

AI Programs at IMD

PwC Switzerland in collaboration with IMD

IMD / Real learning Real impact