

BROOKFIELD, VISA AND JD.COM
SHARE THEIR AI EXPERIENCE
AND INSIGHTS

MOVING FROM THESIS TO
DEPLOYMENT – FROM AI MODELS
TO SECURITY

HOW AI IS REWIRING THE WAY
WE WORK – FROM THE
BOARDROOM DOWN

ENTERPRISE AI

IMPACT

AN FII INSTITUTE PUBLICATION

**TRANSFORMING
BUSINESS**
AI IN THE
ENTERPRISE

FII INSTITUTE


Future Investment Initiative Institute

Impact
on Humanity

FACTS AND FIGURES


HOW IS AI BEING ADOPTED IN THE ENTERPRISE, AND WHAT ARE THE OPPORTUNITIES?

EXTENT OF AGENTIC AI USAGE (%)



[SOURCE: DELOITTE - STATE OF AI IN THE ENTERPRISE, THE UNTAPPED EDGE, 2026]

AI RISKS MOST CONCERNING (PERCENT)



[SOURCE: DELOITTE - STATE OF AI IN THE ENTERPRISE, THE UNTAPPED EDGE, 2026]

\$400 billion

estimated global enterprise AI spending in 2026

[SOURCE: IDC, 2026]

60%

of enterprise generative AI investments in 2024 came from innovation budgets

40%

of enterprise generative AI investments came from permanent or redirected budgets

[SOURCE: MENLO VC, THE STATE OF GENERATIVE AI IN THE ENTERPRISE, 2024]

TABLE OF CONTENTS

- 04 THE NEXT INDUSTRIAL REVOLUTION
- 05 THE HISTORY OF AI IN THE ENTERPRISE
- 10 ACCEPTING THE AGENTIC LEAP
- 14 AN ONGOING CONVERSATION

BACKING AI

- 17 FROM THESIS TO DEPLOYMENT

- 20 REWIRING HOW WORK GETS DONE
- 24 TURNING PILOTS INTO PORTFOLIO VALUE
- 27 A GROWING, BUT UNEVEN, PICTURE

BUILDING AI

- 30 THE NEW ENTERPRISE STACK
- 33 BUILDING TRUST INTO THE TRANSACTION

IMPLEMENTING AI

- 41 WHERE'S MOVING FASTEST ON AI?
- 43 GOVERNANCE, TRUST AND THE AI-READY ENTERPRISE
- 47 HOW CAN FII INSTITUTE HELP?
- 49 KEY TAKEAWAYS TO TRACK



CONTACT US:
info@fii-institute.org

X @FII_Institute1
 in FII Institute
 Future Investment Initiative



AS WE PUBLISH THIS REPORT, it feels like we are standing at the threshold of what many are calling the fourth industrial revolution, driven not by steam, electricity or computing, but by artificial intelligence. But this revolution is different because it's fundamentally reimagining how enterprises make decisions, serve customers and create economic value.

For the last three-and-a-half years, we've talked about AI with wonder. But AI is no longer a futuristic concept or experimental technology. It is operational reality in Fortune 500 companies, financial institutions and businesses across every sector of the global economy.

At FII Institute, we see this transformation firsthand. Through our ACT pillar, we invest in start-ups building the infrastructure and applications that power enterprise AI. Through our FII PRIORITY Compass, which surveys over 60,000 citizens across more than 32 countries, we track how AI adoption intersects with human concerns about work, privacy and economic opportunity.

The data reveals something important: AI sovereignty is emerging as a priority for nations and enterprises alike. The demand for AI that aligns with local values, operates within secure infrastructure and serves specific economic needs is growing.

So we have explored, in this report, Enterprise AI from three different perspectives. First, the macro view: how institutional investors are deploying capital and what adoption patterns reveal about the future. Second, the builder's view: how innovative start-ups are creating agentic AI systems that fundamentally rewire how departments operate. And third, the practitioner's view: how leading enterprises are implementing AI at scale today.

And to do so we have asked for insights from members of the FII Institute community. We are honored to feature contributions from strategic partners who are actively shaping this transformation.

Several things have become clear: the divide is widening between organizations leading the AI charge and those lagging behind. Enterprises that invest now in understanding, deploying and scaling AI will pull ahead. Those that wait will find it increasingly difficult to catch up. Yet it's not just about competition. We need to build AI systems that serve humanity – that create prosperity, expand access to capability and solve real problems. We should not lose track of the fact that AI is the means to those ends, and not an end in itself.

We invite you to engage with these ideas, to question assumptions and to join the debate. Enterprise AI is here. The question is not whether it will transform your business, but how you will harness it to serve your stakeholders and society.

H.R.H. Dr. Maha Bint Mishari Al Saud
CEO, FII Institute





THE HISTORY OF AI IN THE ENTERPRISE

AI has evolved from academic theory into a core enterprise capability, as advances in algorithms, data and cloud infrastructure make machine-driven decisioning practical for business problems. Here are some key steps on the way through that long history.





1950s–1960s

Concept to lab systems

Foundations for enterprise AI began with theoretical work, such as Alan Turing's test and early symbolic approaches that showed machines could mimic parts of human reasoning. While these decades produced little commercial deployment, research on logic, search and knowledge representation created the rule-based expert system architectures enterprises later used for decision support in finance, manufacturing and diagnostics.

PHOTOS: WALTER NURNBERG/SSPL/GETTY IMAGES, PR

1970s–1980s


Expert systems in business

Commercial uptake began when expert systems encoded domain rules for business tasks, exemplified by Stanford's Mycin medical expert system research and Digital Equipment Corporation's RI (XCON) that automated product configuration, saving DEC millions and demonstrating enterprise ROI from encoded knowledge. This era taught companies that AI could reduce human labor on repetitive, decision-heavy work, though systems required heavy maintenance and domain expertise.

RULE-BASED EXPERT SYSTEMS

THE MYCIN EXPERIMENTS OF THE STANFORD
HEURISTIC PROGRAMMING PROJECT

Bruce G. Buchanan
Edward H. Shortliffe





1990s


From rules to statistical methods


The 1990s shifted enterprise AI toward statistical learning and probabilistic models, as computing power and data increased. IBM's Deep Blue showcased specialized search and evaluation, while speech recognition and early data-mining tools began appearing in contact centers and financial services for fraud detection and customer service automation. Businesses moved from brittle rule sets to models trained on data, enabling broader but still narrow commercial applications.

2000s

Big data and predictive analytics

The 2000s saw enterprises collect vast transaction and sensor data, enabling predictive analytics and recommendation systems at scale. Online retailers and banks applied machine learning for scoring, personalization and risk modeling. Combined with improved storage and distributed computing, these capabilities made AI-driven insights a routine part of business intelligence and operations management.





2010s


Deep learning transforms capabilities


Deep learning breakthroughs (convolutional and recurrent networks) and ImageNet-scale datasets produced dramatic gains in vision, speech and language, unlocking enterprise uses in manufacturing inspection, document processing and virtual assistants. Cloud providers and GPUs lowered the barrier to experimentation, prompting a wave of pilots and production projects across retail, healthcare and logistics.

Late 2010s

Cloud platforms and Alops

Cloud AI services and MLOps tooling began making model deployment, monitoring and governance practical for enterprises, turning isolated pilots into repeatable pipelines. This period highlighted how operational maturity such as data pipelines, model versioning and observability are important so organisations could maintain models in production and systematically measure business impact.





2020–2022 Foundation models and NLP at scale

Large language models (LLMs) and foundation models like GPT-3 enabled enterprises to use natural language for search, summarization and code generation, accelerating automation in knowledge work and customer service. Early adopters built chat assistants, document understanding and coding aids, while grappling with hallucination, IP and privacy risks that needed new governance approaches.

2023–2024

Generative AI accelerates adoption

Generative AI's rise in image, audio and text synthesis unlocked an explosion of enterprise use-cases, with a surge in production models and investments in vector databases and retrieval-augmented generation architectures. The early stages saw some missteps and data snafus, which meant enterprises quickly prioritized guardrails, fine-tuning and integration work to control quality and compliance in regulated sectors.





2025

Infrastructure catches up

Enterprise AI moved from model demos to systems that could be run reliably inside businesses. Hyperscalers expanded access to accelerated compute; inference became cheaper; data platforms, vector databases and retrieval systems made company information usable by models; and orchestration tools connected AI to enterprise software. The shift mattered because capability was no longer confined to the model itself. The surrounding stack was becoming mature enough for companies to deploy AI repeatedly, govern it and start pushing it into production.

2026

AI becomes an operating layer

By 2026, those pieces are starting to join up. Models can reason across longer tasks, agents can call tools and act inside software, and enterprises have more mature data, security and governance beneath them. Open-source and proprietary models, public cloud and private deployments increasingly sit inside the same estate. The result is a move from isolated assistants towards AI as an operating layer: something embedded in workflows, measured against business outcomes and capable of coordinating work rather than simply generating content.



ACCEPTING THE AGENTIC LEAP

Moving from chatbots to systems that plan, decide and act.

→ **FOR THE PAST TWO YEARS, ENTERPRISE AI HAS** been dominated by a single interaction pattern: the chatbot. Ask a question, get an answer. Request a summary, receive text. Generate an image, view the result. This pattern of generative AI as a conversational interface has revolutionized how we interact with information. But it is only the beginning.

In recent months we've seen a new paradigm emerging: agentic AI. These systems take action as readily as they generate content. They can plan multistep workflows, make decisions within defined parameters, interact with other systems and autonomously complete complex tasks. While ChatGPT can help you write an email, an agentic AI can draft the email, find the recipient's contact information, send it, follow up if there's no response and update your CRM system accordingly.

That transition is reshaping what's possible inside the enterprise. The organizations that understand this shift first will build significant competitive advantages.

WHY NOW?

Three technical advances are enabling the shift from assistive to agentic AI. First, foundation models have become dramatically more capable at reasoning and planning. The latest models can break down complex tasks into sub-steps, identify dependencies and execute systematically, rather than generating a single response and waiting for the next prompt. Planning is what has previously separated humans from AI – but not for much longer.



We've seen a step-function improvement in model reasoning capabilities over the past year. Models that previously struggled with multistep reasoning can now handle complex planning tasks, maintain context across extended interactions and even revise their plans when new information emerges.

Another contributor is that tool frameworks have also emerged allowing AI models to interact with external systems safely and reliably. APIs, databases and enterprise software can now be called by AI agents within defined guardrails. They handle the complex mechanics of tool selection, error handling and result validation. This is the infrastructure layer that makes agentic systems possible.

Third, enterprises have begun building the data infrastructure and governance frameworks necessary to deploy AI systems that can act on their behalf. You can't have agents making decisions if you don't have trust, transparency and oversight mechanisms in place.

UNLOCKING NEW OPPORTUNITIES

Agentic AI changes the economics of automation in ways that chatbots never could. Take customer service, for instance: an agentic system can handle complex problem resolution, whether it's accessing account data, checking system status, processing refunds or escalating to humans only when necessary. The system can coordinate across departments, pulling information from finance, operations and legal to resolve issues that traditionally required human intermediaries managing handoffs.

A customer with a billing dispute doesn't need to be transferred from billing to collections to legal. An agentic system can pull account history, check relevant policies, calculate appropriate adjustments, process refunds within authority limits and escalate for approval only when necessary.

It's being brought to bear in finance too, where AI can monitor transactions in real-time, flag anomalies, investigate potential fraud by cross-referencing multiple data sources and even initiate remediation actions without human intervention.


Agentic AI in that space can change how businesses operate. Instead of monthly closes and quarterly reviews, finance teams have real-time visibility into financial performance and can respond to issues as they emerge rather than weeks after the fact.

REWRITING SUPPLY CHAINS

In operations, agentic AI systems can optimize supply chains by continuously monitoring inventory, predicting demand, automatically adjusting orders and logistics and coordinating with suppliers – something that could come in handy when disruption occurs, whether that's down to the weather, a port closure or a supplier failure.

AI can change the game from reactive firefighting to proactive tweaking and optimization. Instead of responding to disruptions after they occur, operations teams have systems that can predict disruptions before they happen and automatically adjust to minimize impact.

Those are just a handful of ways in which agentic AI can transform how enterprises work. The common thread is that these systems do more than assist human work. They can execute end-to-end workflows that previously required human coordination across multiple systems. →



→ CHANGING HOW COMPANIES OPERATE

This is a structural change to how organizations function. When AI can plan and execute workflows, the boundary between what humans do and what AI does shifts. Routine multistep processes become fully automatable while decision-making becomes distributed across

human-AI collaboration rather than centralized in human managers. Organizational hierarchies get flattened along the way. It's disruptive by name and by nature.

But it also poses questions that leaders need to consider. Which processes should be fully automated? Where do humans need to remain in the loop? How do you maintain accountability when decisions are made by algorithms? How do you manage risk when systems can act autonomously?


The organizations that succeed will be those that approach agentic AI systematically by identifying high-value workflows, designing appropriate guardrails, measuring outcomes rigorously and expanding incrementally as confidence and capability grow.

THE NEW ECONOMICS OF AUTOMATION

What makes this such a trying time is that the economics of agentic AI are fundamentally different from previous waves of automation. We've been through this before, but never in such a way as with agentic AI. Traditional automation required explicit programming for every step, making it expensive and brittle. Agentic AI can figure out the steps itself, adapting to changing conditions without reprogramming.

What that means is automation can now be applied to processes that were previously too complex or variable. Processes that required judgement, coordination across systems or adaptation to changing circumstances can now be handled autonomously.

Systems that were considered "human-only," because they required reasoning or coordination, can now be handled by AI systems. This doesn't mean



TOKENIZED US TREASURIES AUM:

30–50%

acceleration in business processes

25–40%

reduction in low-value work time for humans

60%reduction in risk events in
finance/risk monitoring pilotsSOURCE: BCG: HOW AGENTIC AI IS
TRANSFORMING ENTERPRISE PLATFORMS

humans become obsolete – but instead, humans shift to higher value activities like strategy, creativity and relationship-building. For enterprises that can automate more of their operations the change is profound. They can get lower costs, faster cycle times and better consistency. And more importantly, they can free their human talent to focus on work that creates genuine differentiation.

FROM PILOT TO PRODUCTION

The enterprises seeing the earliest success with agentic AI are those that start with well-defined, high-value processes and expand from there. They begin with processes where the rules are clear, the systems are integrated and the risk is low. They prove the technology works, build trust and then slowly expand to more complex use cases.

An incremental approach means organizations can understand where guardrails are needed, where human oversight is essential and how best to measure ROI. Starting small, proving value, then expanding aggressively is the model to follow.

Boiling the ocean by deploying agentic AI across the enterprise simultaneously doesn't work. The organizations that succeed are those that methodically build capability and confidence, and scale over time.

The agentic leap is here. The technology exists, the infrastructure is maturing and early adopters are already seeing results. The only question is how quickly your organization will adapt. ■



AN ONGOING CONVERSATION

AI isn't just a flash in the pan – nor is it something new to our FII Institute community, as these highlights from our past events show.



"Just about every company by the fourth quarter will start using agentic AI, whether in the back office through a procurement assistant or in HR through a benefits assistant."

BETSY ATKINS

Google Cloud Advisory, Board; Member of Board of Directors, GoPuff & Wynn Las Vegas, speaking at FII PRIORITY Miami



"We always start by ensuring the enterprise backbone is solid, because if you do not have a strong foundation, you simply cannot build a scalable AI driven organization on top of it."

GARY COHN, Vice Chairman, IBM, speaking at FII PRIORITY Miami

"AI enables mass customization at scale at much lower cost, and ultimately those companies that can harness these technologies to build sustainable advantage and compound value over time are the ones that will endure."

YIE-HSIN HUNG

President & CEO, State Street Investment Management, speaking at FII PRIORITY Miami



“AI will drive both efficiency and creativity across the industry allowing teams work faster, tell better stories and ultimately transforming how content is created, distributed and experienced by audiences worldwide.”

DAVID ELLISON

Chairman & CEO, Paramount Skydance,
speaking at FII PRIORITY Miami



“In NEOM, AI is not being used as a support function but as a core decision-making layer, deciding where electrons are more efficiently directed across the energy and water systems.”

RAYAN FAYEZ

Deputy CEO, NEOM, speaking
at FII PRIORITY Miami

“Using AI to optimize energy usage on rigs, including shifting from engines to battery systems under certain loads, has helped reduce emissions by nearly 20% while improving overall environmental sustainability.”

DR. MOHAMED FAROUK

Vice Chairman & Group CEO, ADES
Holding, speaking at FII PRIORITY Miami



BACKING

AI

The capital markets have spoken: AI is the investment theme of the decade. From public equity markets to private equity, from venture capital to sovereign wealth funds, institutional investors are deploying capital at unprecedented scale to capture the AI opportunity.

FROM THESIS TO DEPLOYMENT

Here's how institutional allocators are using AI – and investing in it.



FOR INSTITUTIONAL ALLOCATORS OF ALL

stripes, from public-market managers and private-equity sponsors to sovereign funds, AI has become both an investment opportunity and an operational tool. The firms that will thrive in the coming decade are those that master both investing in AI companies and using AI inside their own operations to sharpen research, diligence and risk.

Ask around the investment space and the biggest secret is how AI is becoming a source of alpha. The investment thesis is straightforward: AI will reshape every industry, and the companies that adopt AI effectively will outperform those that don't. For allocators, this creates both an opportunity and an imperative.

The companies building AI infrastructure, models and applications are generating extraordinary returns. Early investors in these companies have captured outsized gains. As AI adoption accelerates, the opportunity set is expanding from pure-play AI companies to traditional businesses deploying AI effectively.

With that in mind, portfolio companies that adopt AI will outperform those that don't. Private equity (PE) firms have a responsibility to drive value creation in their portfolio companies, and AI is one of the most powerful value creation levers available to them.

That dual reality of AI as both an investment opportunity and an operational tool is reshaping how allocators think about their businesses. The most sophisticated firms are treating AI as a core capability, building internal expertise while deploying capital to the most promising AI companies. →



→ OPERATION ALPHA

Beyond the investment thesis, AI is transforming the way that allocators do their work. In public markets, hedge funds and mutual funds are using AI to analyze alternative data sources, identify trading signals and manage risk.

Given AI's power at parsing through big data, AI is being used in all sorts of ways. Satellite imagery analysis reveals economic activity before earnings reports. Earnings call sentiment analysis can identify management confidence beyond the prepared remarks, while supply chain monitoring can head off any revenue surprises. Using AI effectively is a bit like a cheat code, generating insights that were previously inaccessible.

And in private equity, firms are using AI to improve due diligence, assess operational efficiency in target companies and identify value creation opportunities post-acquisition. AI can analyze thousands of documents in seconds, identify cost optimization opportunities, and benchmark company performance against peers. Something that used to take weeks can now be accomplished in days.

That trickles up to sovereign wealth funds (SWFs), too, with them using AI to model long-term economic scenarios, assess strategic risk and make allocation decisions across asset classes.

AI can simulate thousands of scenarios and identify

portfolio optimizations that human analysts might miss. When talking about the long-term investments that are typical of SWFs, with time horizons spanning decades, that can make a meaningful difference.

A POWERFUL FEEDBACK LOOP

It all adds up to the creation of a powerful feedback loop: as allocators get better at using AI, they make better AI investments. Then as those investments perform, allocators have more capital to deploy.

For a PE firm that uses AI to identify which portfolio companies are ready for AI adoption, then helps those companies to deploy AI effectively, AI can result in improved performance and higher exit multiples. The selfsame AI capabilities that helped to identify the opportunity can be redeployed to find the next one – and repeat the pattern over and over.

If there's one thing that we've seen throughout successive discussions at repeated FII Institute events around the world, it's that those firms that invest early in AI capabilities see immediate returns, which fund further investment.

Meanwhile, the firms that wait find themselves falling further behind with each passing quarter.

BUILD, BUY OR PARTNER?

For allocators, as well as for the companies they back, the practical question is how much AI capability to own.

Building makes sense where proprietary data, workflows or decision-making are a source of advantage and the organization has the talent to maintain what it creates. Buying is often faster for common functions where mature tools already exist and differentiation comes from adoption rather than the soft- →



\$66 billion

sovereign-owned investment into
AI and digitalization in 2025

(SOURCE: GLOBAL SWF REPORT 2026)

88%

of asset allocators plan to grow
AI allocations in the next year

75%

of asset allocators planned to
grow AI allocations in 2024

(SOURCE: DYNAMO SOFTWARE SURVEY, 2025)

19.95%

EU enterprise AI adoption in 2025

55%

of which are large firms

17%

of which are small firms

(SOURCE: ALICE LABS AI ADOPTION INDEX, 2026)

→ ware itself. Partnership sits between the two: firms can keep control of data, governance and business logic while leaning on specialist providers for models, infrastructure or integration.

The choice is rarely permanent. Many organisations will buy to prove value, partner to integrate systems and build selectively where AI becomes strategically important. The risk is treating every use case the same: overbuilding commodity capability can waste capital, while outsourcing the parts that make a business distinctive can create dependency. The investment case therefore extends beyond who owns the best model to who makes the clearest choices about where AI belongs.

THE GREAT SHIFT

Perhaps most importantly, AI is shifting from a cost line to a profit center for capital allocators. The firms that treat AI as a strategic capability rather than as a support function, AI's boosters believe, will generate better investment returns, operate more efficiently, and build sustainable competitive advantages.

But there's no doubt that shift requires rethinking how AI is organized and funded within the firm. Rather than treating AI as an IT initiative or experimental project, firms leading the way are building centralized AI capabilities that serve the entire organization, and hiring AI talent not just for technology teams but for investment teams, risk teams and operations teams. The most advanced are creating career paths for AI professionals within the investment organization.

They're often the ones capturing the full value of AI – both as an investment theme and as an operational tool. Those that treat AI as a peripheral capability will find themselves competing against firms that have made AI central to everything they do.

And increasingly, those firms realize the question is no longer whether allocators should invest in AI. It's how quickly they can build the internal capabilities to use AI effectively while deploying capital to the most promising AI companies. As we've found from the FII Institute's community, the firms that master moving from thesis to deployment will be the winners in the AI era. ■

REWIRING HOW WORK GETS DONE

Agentic workflows are transforming the enterprise, from how services are delivered to who is employed.



BEYOND THE SURFACE-LEVEL ADOPTION OF

chatbots answering questions and copilots drafting emails, a deeper transformation is underway. Agentic AI workflows are fundamentally rewiring how departments operate, make decisions and interact with one another.

It requires business processes to be rewired from the ground up, with AI as an active participant rather than a passive tool.

Advisory company Gartner has projected that by 2028, roughly one-third of enterprise software applications will include agentic AI capabilities, up from less than 1% in 2024. Gartner also suggested that around 15% of day-to-day work decisions could be made autonomously by such systems within that timeframe.

But why is this happening, and how is it rewiring how work is done?

SLOWLY AT FIRST, THEN ALL AT ONCE

Frontier models crossed a reasoning threshold in 2024 and 2025 that made multistep planning reliable enough for production use. Instead of one-shot answers, models could now decompose a goal, attempt a step, observe the result and adjust – the basic loop on which any agent depends.

At the same time, the plumbing that makes AI work became standardized. Anthropic's Model Context Protocol, released in late 2024, gave agents a common way to connect to enterprise tools, data sources and applications without bespoke integration work for each one. Comparable initiatives followed across the industry. The result was that an agent built on one platform could increasingly reach into a company's CRM, ticketing system, data warehouse and email without months of custom engineering.

Compute economics also improved. The cost of frontier-class inference (delivering AI responses to queries) fell by roughly an order of magnitude between 2023 and 2025,

making it economically viable to run agents that consume far more tokens than a chatbot, because they reason, retry and verify things before acting.

Every major enterprise software vendor shipped an agent platform. Salesforce launched Agentforce. Microsoft pushed Copilot Studio and a catalogue of pre-built agents inside Microsoft 365. ServiceNow, SAP, Oracle, Workday and Google Cloud all introduced agent-building frameworks. FII Institute also announced its own agentic AI tool specifically designed for governments, SAGE, at FII9 in October 2025. The decision facing most CIOs is no longer whether to deploy agentic AI, but on whose stack.

WHAT AI HELPS UNLOCK

While adoption is still early, the widespread use of AI means we have early examples of success. Klarna disclosed in 2024 that its AI assistant was performing work equivalent to 700 full-time agents within months of launch, handling two-thirds of customer service chats with resolution times falling from 11 minutes to under 2. Banks, telcos and airlines have since published similar improvements in their own areas of expertise.

In software engineering, agents have moved from autocomplete to autonomous task completion. Coding agents now write code, run tests, open pull requests and autonomously respond to review comments. GitHub reported that developers using its Copilot tools complete certain tasks 55% faster; agentic coding tools released in 2025 widened that gap considerably for well-scoped work.

In finance and operations, agents are being deployed against the long tail of administrative work that resists traditional automation: reconciling invoices that don't quite match, chasing vendors for missing tax forms, preparing month-end variance commentary, triaging procurement requests against policy. McKinsey's State of AI research has consistently shown finance and supply chain among the functions where generative AI delivers measurable cost reduction, and agentic deployments →



→extend that further by closing the loop from analysis to action.

In sales, several Fortune 500 firms have publicly described pilots in which a single seller is now supported by a team of specialized agents handling research, follow-up and administrative load.

HOW AI REWIRES ORGANIZATIONS

When agents can reliably execute multistep workflows, the unit of work inside a company starts to change. Processes that were designed around human handoffs collapse – drastically.

That has consequences for organizational design. But not necessarily in the way that people fear: while much has been made of the International Monetary Fund's warnings that 40% of jobs will be affected in some way by AI, "affected" doesn't mean lost. It just means changing.

New roles have emerged in response to the rise of AI, like agent supervisors, evaluation engineers, prompt and policy designers, and more broadly people responsible for deciding which decisions a machine is allowed to make.

Risk management also gets harder for enterprises in an AI-enabled world. An agent that can act has to be governed differently from a chatbot that can only suggest. Companies are now writing policies on which actions require human approval, how agent decisions are logged, how errors are caught and what happens when an agent makes a confident mistake at scale.

Most industry analysts have projected that a majority of enterprises using generative AI will have launched agentic pilots by next year – suggesting that those who don't risk losing out.

MORE HASTE, LESS SPEED

With all that in mind, and listening to the conversations at FII Institute meetings around the world, it's hard not to see the way work gets done permanently changing. And yet it's worth exercising some caution – tapping the brakes, rather than slamming them on, or coasting without them. Agentic AI isn't a finished technology yet. AI models still hallucinate. And while the fully autonomous enterprise agent that handles open-ended objectives across systems with no supervision is closer than it was, it's still not here.

But what is here is a phase change. The conversation about AI inside the enterprise has shifted from "can it answer" to "can it act."

And with that, executives are asking less whether AI – agentic or not – will change how work gets done. Instead they're asking the tougher questions of which processes to rewire first, which decisions to delegate – and beyond that, how quickly the organization can adapt to a workforce in which not all their workers are necessarily human.



JOIN THE GROWING FII INSTITUTE COMMUNITY

Whether you're an organization trying to make AI work within your enterprise, a startup or big business supplying AI tools to global clients, or a government representative trying to figure out how to capitalize on the potential of the technology while keeping business as usual running, the FII Institute community can support you. We are something more than a network – we're a global force for action.

With regular PRIORITY events across the globe, alongside our flagship event, celebrating its 10th edition this year in Riyadh, there are plenty of opportunities to congregate with likeminded people and share best practices, as well as debate the key topics that will shape the future.

That's supplemented by Impact Reports such as these, which tackle the key issues and set the scenes for our regular gatherings. Besides that, regular insights from members and key Strategic Partners provide further ideas of how to tackle the issues of our age in frequently published op-eds.

As a member of FII Institute, you join The 1500 – a select community of executives, researchers, and entrepreneurs focused on building solutions for tomorrow's biggest challenges. Members collaborate on high-impact initiatives, contribute to global dialogues, and gain access to exclusive platforms including FII Conference and FII Global Priority Summits.

With members from all four corners of the globe, and representing every aspect of business, politics and NGOs, we offer a comprehensive oversight of what's changing the world, right now.

So join us in our growing FII Institute community



TURNING PILOTS INTO PORTFOLIO VALUE

Brookfield's approach to enterprise AI combines investment discipline with hands-on implementation.

→ **ORGANIZATIONS AROUND THE WORLD ARE** putting plenty of AI pilots into the field, testing whether the technology can improve their processes and bolster their bottom line. But for an investor, the big question is when and how to turn those pilots into production. Who will do so? What will improve? And how will that improvement be measured? Brookfield's approach to enterprise AI begins with the work required to answer those questions.


The firm views enterprise AI like any other operational change that could make an investment more valuable: can it increase revenue, bring down costs, lead to new products or leave a company in a stronger competitive position? It then weighs that against the cost and difficulty of putting it into practice, the risks involved – and whether it can work at scale.

Core to all this is the nitty-gritty of making it work on the ground. “We believe a credible AI strategy is defined by execution,” Brookfield says. It's most interested in use cases that have made it out of the experimental stage

and into production, with results that can be linked to financial performance. And it's not picky about size, either – it recognizes that several smaller projects with numbers behind them can be more meaningful than one grand plan for what AI might one day achieve.

LOOKING INSIDE THE BUSINESS

Even the most promising projects require the right organizational infrastructure to deliver meaningful change. Which is why Brookfield also looks at the organization built up around an AI project. A company may have built a good tool, but does anyone have →



→ responsibility for it? Is there a team behind it, a clear reason for using it and a place for it in the work people already do? Those questions matter to Brookfield, because if nobody owns the process technology is supposed to improve, or workers have to step outside their usual systems to use it, a successful pilot can stall and never get out of first gear.

Brookfield applies the same thinking internally. The firm uses AI in sourcing, due diligence and portfolio management, including tools that surface operational insights and compare performance across companies. Its investment professionals are also building applications that support their workflows, from early screening through preparation for investment committee decisions.

Trusting development to the hands of people who already understand the job can certainly speed things up and produce more effective tools. But without due care and attention, it can also create another problem: a proliferation of tools and applications that may not integrate effectively or deliver clearly measurable impact. A key part of Brookfield's job is therefore to decide what is actually useful, what can be picked up elsewhere and what needs more work before it can be relied upon to be rolled out across an organization.

EXECUTE AND ITERATE

Part of the answer has been to create an AI Value Creation Office, staffed by more than 30 people drawn from investment, operations and technology. It works across Brookfield's investment platforms.

A shared repository helps internal teams examine what has worked elsewhere rather than begin every initiative from scratch. The point is simple: if another company has already spent time finding out what works – or what doesn't – then there's little sense in starting from zero. Those choices are then prioritized by quantifying the expected impact before committing resources, while regular meetings allow platform leaders to compare progress and identify opportunities across the portfolio.

Those regular catch-ups are important, as Brookfield has more than 800 identified AI use cases at different stages of execution. The meetings allow the firm to see where projects lose momentum or don't quite work, and to direct support to overcome the hurdles stopping an idea from achieving what it's meant to.

"The biggest lesson is that the gap between having an idea and getting it into production is where most companies lose momentum," the firm says.

From that, Brookfield has learned that progress tends to be faster when operating leaders are →



→ pushing the work forward, there's a business case from the outset, and the technology is being used to fix a problem employees already recognize in their daily work. Senior backing matters, too. But clear ownership remains critical: an individual or team still has to be accountable for implementation.

"Companies without a clear owner for the use case often see it drift," Brookfield said. Instead of treating AI as a separate program, embedding it in the way a company intends to create value gives managers a reason to overcome the procedural and organizational questions that often plague rollout.

Sharing any experiences across the portfolio is another advantage. When one part of the business demonstrates that an approach works, others can learn from that.

RESULTS NEED A BASELINE

Using AI doesn't matter if its results can't be quantified. Brookfield measures its returns through a raft of measures including cost savings, revenue potential, time saved and improvements in process efficiency. Applying the same framework across businesses helps the firm compare opportunities and examine where the gains are emerging.

And those gains can be significant. In manufacturing and supply chains, cumulative savings of millions have been achieved through optimization. In media and data businesses, Brookfield has found projected annual savings in the tens of millions, as well as work that once took months being completed in days.

And in customer operations, the firm has seen a 40% reduction in hiring costs per full-time equivalent and a fivefold increase in hiring speed, alongside tools supporting customer interactions.

CREATING KEY INFRASTRUCTURE

For Brookfield, the AI investment story goes well beyond software. The firm sees data centers, power and computing capacity as increasingly connected assets. And as AI demand grows, reliable energy and the infrastructure that links supply to consumption become central to the conundrum.

That last part is becoming an increasingly important constraint. Grid access, electricity generation and transmission connections remain constrained in many markets. A company can have a ready-to-run AI product, but still find itself stalled because the power or data center capacity it needs doesn't exist. And unlike software, those problems cannot usually be fixed with another update pushed out via the internet. The physical infrastructure that makes this stuff happen, from power stations to data centers, take time and vast sums of capital to build.

Brookfield believes it's vital to build out the infrastructure supporting AI, backed by long-term contracts and adaptable assets.


Across both portfolio operations and infrastructure, the principles are the same: find the constraint preventing value from being created, assign responsibility and invest in a way to remove it. Enterprise AI can play an important role in doing so. But, as Brookfield's experience suggests, realizing its potential depends on the infrastructure, ownership and execution around it.



A GROWING, BUT UNEVEN PICTURE

How AI adoption is occurring across enterprises depends on where you are in the world – and how big you are.

EXTENT OF AGENTIC AI USAGE (%)



SOURCE: FII PRIORITY COMPASS 2025

BUYING BIG INTO AI

\$344.6b

global private investment in AI in 2025


127.5%

increase year-on-year

SOURCE: STANFORD AI INDEX, 2026

GLOBAL PRIVATE INVESTMENT IN AI BY GEOGRAPHIC AREA, 2025

Total investment, in billions of US dollars



SOURCE: THE STATE OF ART IN THE ENTERPRISE, DELOITTE, 2026

HOW IMPORTANT ARE THE FOLLOWING DIGITAL SERVICES IN YOUR DAILY LIFE?

■ ALL ■ GLOBAL SOUTH ■ GLOBAL NORTH

ACCESS TO HIGH-SPEED INTERNET



FINANCIAL SERVICES (E.G. ONLINE BANKING, INVESTING)



E-COMMERCE AND ONLINE PAYMENTS



DIGITAL TOOLS FOR E-GOVERNMENT SERVICES



DIGITAL PLATFORMS FOR SOCIAL INTERACTION



ONLINE LEARNING OR SKILL DEVELOPMENT COURSES



DIGITAL TOOLS OR APPS TO MANAGE YOUR HEALTH OR WELLNESS



REMOTE WORK TOOL



GENERATIVE AI TOOLS (E.G. CHATGPT, MIDJOURNEY)




SOURCE: FII PRIORITY COMPASS 2025

LEVEL OF PREPAREDNESS FOR AI ADOPTION

in %

■ NOT PREPARED ■ SLIGHTLY PREPARED ■ MODERATELY PREPARED ■ HIGHLY PREPARED



SOURCE: THE STATE OF ART IN THE ENTERPRISE, DELOITTE, 2026

BUILDING AI

The enterprise AI start-up ecosystem is booming. From foundation model builders to vertical application specialists, from data infrastructure to developer tools, thousands of start-ups are racing to capture the AI opportunity.




PHOTO: LOPODOKOFF/GETTY IMAGES/ISTOCKPHOTO

THE NEW ENTERPRISE STACK

Mapping out the emerging enterprise AI technology landscape is the best way to understand how to navigate it.

→ **BENEATH EVERY AI APPLICATION, WHETHER IT'S** a chatbot, agent or predictive model, is a technology stack, based on AI, that has embedded itself into organizations within a matter of months. Understanding the stack is essential for anyone making decisions about enterprise AI adoption.

The reason? The choices you make at each layer determine what's possible, what it costs and how much control you have over your AI destiny.

LAYER 1: FOUNDATION MODELS

At the base of the stack are the foundation models – large language models and multimodal AI systems that provide core intelligence capabilities.

OpenAI, Anthropic, Google and Meta have emerged as the leaders in proprietary foundation models. Llama, Mistral, and Falcon have emerged as the leading open-source alternatives. The competition is fierce, with capabilities advancing on a weekly basis.

For enterprises, the choice at this layer dictates everything that follows: proprietary models offer cutting-edge capabilities, but create dependency on external providers. Open-source models offer more control, but require more expertise to deploy and maintain. Most enterprises will ultimately use both, going for proprietary models for applications where capability is paramount, open-source for applications where control or cost is the priority.

The economics of foundation models are also evolving. Training costs are measured in hundreds of millions of dollars, creating a moat around the largest players. Because of that, large AI providers are hiking prices, pushing users to local-hosted LLMs. →

→ LAYER 2: ORCHESTRATION AND TOOLS

Sitting above the foundation models is the orchestration layer – the frameworks and tools that make it possible to build applications on top of foundation models.

Here, LangChain, LlamaIndex and Microsoft Semantic Kernel provide developer frameworks for AI applications. Vector databases like Pinecone and Weaviate enable retrieval-augmented generation, allowing AI systems to access enterprise data. MLOps platforms like Weights & Biases and MLflow provide monitoring and governance.

This layer is where most enterprise development happens today. It's the one where developers connect foundation models to enterprise data, build agentic workflows and implement the business logic that makes AI useful. The winners in this layer will be the companies that make it easiest for enterprises to build, deploy, and manage AI applications.

The competitive dynamics in this layer are still uncertain. Consolidation is taking as larger platforms acquire point solutions, while platform vendors embed orchestration capabilities directly into their products.

LAYER 3: VERTICAL COPILOTS

Above the orchestration layer are vertical copilots – AI applications specialized for specific roles or functions.

Consider tools like Harvey, which layers AI into legal work, analyzing contracts and conducting legal research, or AlphaSense, which provides AI for financial research, analyzing earnings calls and SEC filings. Vitals provides AI for healthcare, assisting diagnosis and treatment planning. These tools are trained on industry-specific data and designed to solve industry-specific problems. They understand industry terminology, comply with industry regulations, and integrate with industry systems.

For enterprises, the choice at this layer is between buying vertical copilots off the shelf or building custom solutions. The right answer depends on strategic priorities, data assets, and appetite for technology development.

LAYER 4: AGENTIC WORKFLOW PLATFORMS

The emerging top layer of the stack is agentic workflow platforms – the systems that can plan, execute and coordinate multistep workflows across enterprise systems. HUMAIN One, Salesforce Agentforce, and Microsoft Autogen are leading examples. These platforms go beyond simply answering questions. They take action. They can coordinate across departments, integrate with enterprise software and operate autonomously within defined guardrails.

This is where enterprise AI is heading next. The enterprises that invest now in understanding and experimenting with agentic platforms will be positioned to capture the first-mover advantages as this technology matures. The laggards will find themselves trying to catch up as competitors deploy AI that can execute workflows autonomously.

THE LAYER BENEATH THE STACK

Picking the model is only one part of enterprise AI. Beneath it sits the physical and digital infrastructure that determines how quickly, cheaply and securely AI can run. Hyperscalers and public-cloud platforms provide access to accelerated compute; GPUs and other accelerators do the heavy lifting; inference infrastructure serves models once they are deployed; and enterprise data platforms supply the context that turns a general model into something useful inside a business.

That infrastructure increasingly shapes strategic choices. A company using public cloud can scale quickly, but becomes more dependent on external providers. Private or sovereign deployments offer more control, but require more infrastructure and operational expertise. Data platforms matter just as much: without clean, accessible and well-governed enterprise data, even a capable model has little to act on.

Seen this way, the enterprise AI stack starts below the foundation model. Compute, cloud, inference and data are not background plumbing. They set the cost, control and resilience of everything built above them. ■

GO SOVEREIGN OR PRIVATE?

As enterprises build their AI stacks, a critical question emerges: where does your AI live?

SOVEREIGN AI systems operate within national borders, comply with local regulations, and are controlled by local entities – important for enterprises in regulated industries or countries with data localization requirements. Massive investment in sovereign AI infrastructure globally – particularly in regions that want to reduce dependence on US and Chinese AI providers – is taking place.

PRIVATE AI systems are deployed privately within an enterprise, either on-premise or in a private cloud. This is increasingly important for enterprises with proprietary data or trade secrets that cannot leave their control. Cloud providers are rushing to offer private AI deployment options, and a new generation of AI appliances is making on-premise deployment increasingly viable.

PUBLIC AI systems are accessed via public APIs and hosted by third-party providers. This is the easiest deployment model, but creates dependency on external providers and potential data privacy concerns. Most enterprises begin with public AI for experimentation but move to private or sovereign options as they scale to production.

Most enterprises will ultimately use a hybrid approach: public AI for non-sensitive workloads, private AI for sensitive operations and sovereign AI where regulatory requirements demand it.



BUILDING TRUST INTO THE TRANSACTION

AI is moving from helping consumers discover products to helping them make purchases. As commerce enters the agentic era, the challenge is ensuring that trust, security and consumer control evolve alongside new forms of automation.



AI AGENTS ARE MOVING RAPIDLY FROM HELPING

consumers discover products to making decisions and completing transactions on their behalf. Giving AI access to a small corner of the web is one thing, but allowing it to complete a payment is another challenge entirely.

For Visa, this is not a theoretical challenge. For decades, the company has helped establish trust among consumers, merchants and financial institutions in an increasingly digital world. As commerce enters the agentic era, that same infrastructure must help participants understand who is acting, on whose behalf and under what authority. →



→ “As commerce becomes more AI-driven, the real test of trust will be at the point of payment, where control, security, and confidence have to come together,” Visa says. The challenge for enterprises is to build that confidence into systems that can operate reliably at scale.

AI is already changing consumer behavior. Visa’s Global Travel Intentions Study found that 50% of travelers surveyed used AI to discover travel ideas, 48% used it for reviews and recommendations and 45% used it to find flights and accommodation. As AI becomes embedded in how people discover destinations, plan journeys and make purchases, it can help create more seamless and personalized commerce experiences.

In Saudi Arabia, this opportunity is emerging alongside the Kingdom’s broader investment in digital transformation, tourism and innovation. Visa research with Fast Company Middle East found that 63% of surveyed business leaders in the country expect agentic commerce to disrupt their industry within two years, while 43.3% are evaluating AI-assisted purchasing and 20% are developing pilots. The findings point to a pragmatic market focused on governance, trust and measurable value before scaling.

FROM FRAUD DETECTION TO INTELLIGENT COMMERCE

Visa recognised the transformative power of AI more than 30 years ago, and its importance continues to grow. In 1993, Visa became the first network to deploy AI-based technology for risk and fraud management, pioneering the use of AI models in payments.

That experience has helped Visa develop a deep understanding of how AI can support the payments ecosystem, where decisions must be made quickly, sensitive information protected and suspicious activity distinguished from legitimate purchasing.

Visa has long played a leading role in securing digital payments, investing more than \$13 billion in technology over the past five years, including efforts to reduce fraud and strengthen network security. This investment reflects the broader technology ecosystem required to ensure that AI is deployed with safety and security at its core.

Today, Visa’s technology platform is among the most powerful examples of the tangible benefits of AI. With several hundred AI models in production, powering more than a hundred products, Visa’s AI and deep-learning capabilities help address long-standing challenges and pain points for buyers, sellers and financial institutions.

RECOGNIZING THE AGENT

Agent-led commerce introduces a new participant into

a payments ecosystem built on identity, authorization and accountability. An agent may act for a consumer or business, but does not remove the need to verify who is acting, define what the agent is permitted to do and preserve clear accountability for the transaction. Building trusted agentic commerce requires foundational capabilities spanning identity, authentication, permissions, spending controls and dispute management. Visa Intelligent Commerce supports these principles by helping define what an agent is allowed to execute through spending limits, approval thresholds and permission layers. Consumers must remain able to set boundaries and understand when and why an agent is acting on their behalf. The same principles must extend across the ecosystem. Visa Agentic Ready helps issuers understand and test agent-initiated payments, while →



As commerce becomes more AI-driven, the real test of trust will be at the point of payment.”



ALI BAILOUN

is Senior Vice President and Group General Manager for Visa across Saudi Arabia, Bahrain and Oman, leading the company’s strategy and long-term market plans across these three priority growth markets. He leads Visa’s commercial growth, regulatory engagement, and ecosystem development, working with financial institutions, fintechs, digital banks, government bodies and market infrastructure partners to advance the region’s digital payments landscape.



→ merchants need visibility into whether an interaction is legitimate and whether their digital experiences are ready for agents to navigate.

Visa Agent Score and Visa's Agentic Directory are examples of capabilities designed to give merchants greater visibility into how agents navigate their websites, who those agents represent and whether they are able to transact. Together, these tools support the broader trust architecture needed for agentic commerce to scale.

PROTECTING PAYMENTS

Security must evolve as quickly as the threats. While cybercriminals use AI to increase the speed, scale and sophistication of attacks, payment networks must distinguish legitimate activity from fraud in real time, without unnecessary friction for consumers and businesses. AI fraud detection uses continuously learning artificial-intelligence and machine-learning models to identify unusual behavior across large volumes of payment data. By analysing transaction details, device fingerprints and behavioral signals in real time, these systems can generate risk scores and help approve, challenge or decline transactions within milliseconds.

This approach can strengthen security while supporting a smoother customer experience. Recognizing legitimate purchasing patterns helps reduce false declines, while identifying behavioral anomalies can help detect account takeover and other suspicious activity. Visa's launch of the Visa Threat Intelligence Platform forms part of its broader approach to helping financial institutions identify and respond to cyber threats before they result in fraud or financial loss.

WHAT'S COMING NEXT

If AI is to make economies more human, not less, businesses must ensure that AI-enabled interactions remain secure, transparent and accountable. Fully autonomous shopping is not yet operating at scale, making human oversight, trusted infrastructure and clear permissions essential as adoption develops.

It's also a challenge for the merchants, banks and technology providers that control individual parts of the broader environment. Progress depends on their systems working together – and heading off risks before they arise.

Trust in payments is more than identity checks or risk decision-making. It also requires clear accountability and the ability to respond when something goes wrong. Used responsibly, AI can help people discover opportunities, access services, travel more easily and participate more fully in economic life. The question is not whether AI will transform economies, but how societies can harness that transformation responsibly. Countries and organizations that build secure, transparent and resilient digital foundations today will be best positioned to realize AI's economic potential tomorrow.

MAKING INTELLIGENCE OPERATIONAL

Sandy Xu, CEO, JD.com, explains why connecting decisions, physical operations and people matters more than collecting impressive demonstrations.

SANDY XU

is the Chief Executive Officer of JD.com, also known as JINGDONG. JD.com is a leading supply chain-based technology and service provider that opens its technology and infrastructure to partners and industries to drive productivity and innovation. JD.com is listed on Nasdaq (NASDAQ: JD) and the Hong Kong Stock Exchange (HKEX: 9618). It ranks No. 41 on the 2026 Fortune Global 500.



TO GET A PARCEL MOVING THROUGH A LOGISTICS

network, you need to have a series of decisions operating as one. Stock must be positioned, goods picked and sorted, transport assigned and the final journey planned. Tinkering around the edges – such as making one decision faster – can help with the process. But connecting those decisions can change how the whole operation works – for the better.

That's the distinction **Sandy Xu**, CEO, JD.com, draws between AI experimentation and enterprise capability. The company's starting point is the same question it asks of other investments in its business. Does the technology improve the customer experience, lower costs and increase efficiency?

"It becomes an enterprise capability only when it is embedded in real business processes and consistently delivers measurable results," Xu says. JD.com has applied AI across more than 3,000 scenarios in businesses spanning retail, logistics, healthcare, industrial supply chains and food delivery. But as every major organization is finding with their AI adoption, the challenge is to connect that activity to useful operational change.

WHERE DECISIONS BECOME ACTIONS

Xu sees particular potential where information meets the physical world. Language models already help JD.com employees work with information in functions like customer service and finance. It's a little trickier in logistics and industrial operations, which add equipment, inventory and people to the equation.

"AI often creates the greatest value not where there is the most information, but where information can ultimately be translated into decisions and actions," she says. A model's understanding becomes commercially relevant when it helps determine what to buy, where to put it or how to fulfil a customer's request.

In logistics, JINGDONG Logistics' MetaBrain connects AI with planning and decision-making across the network. Xu said it has been used in more than 1,000 →

→ logistics and supply chain scenarios, with the objective of moving from improving individual processes towards forecasting and coordination across connected operations.

AI can also be used in physical environments thanks to increasing innovation in embodied AI, in relatively controlled environments, including warehouse navigation and moving goods – with Xu expecting its role to expand further as perception, reasoning and execution improve.

LEARNING THE INDUSTRY

AI excels at general knowledge, but industrial procurement shows we also need more specialized AI capabilities built for specific industries and use cases. Products have detailed specifications, and data standards can differ depending on which business or supplier you're working with. So JINGDONG Industrials has developed JoyIndustrial, an industrial supply chain foundation model combining AI with accumulated product and supply chain knowledge. Xu says the business used more than 70 AI agents across its value chain in the first half of 2026, including applications in product data, matching supply and demand, procurement and fulfilment.

Yet the larger opportunity, Xu argues, lies in joining work that has previously been fragmented. Marketing content can be generated more efficiently within one department, while it's possible now to use AI to connect customer demand to product selection, inventory forecasting and fulfilment. All that of course requires changes within the company itself, too, from responsibilities to information flows.

Xu draws a parallel between AI and electrification: businesses gained more when they redesigned production around a new capability than when they simply substituted equipment within an existing process. Her argument for AI suggests that organizations need to examine what work should become once information and decisions are able to move differently.



MEASURING THE DIFFERENCE

At JD.com, there has been a gradual process of connecting and upgrading operations, bringing individual workflows into a more integrated whole, says Xu. Leaders have to decide who owns a process, how teams cooperate and what governance applies when an agent participates in work spanning several functions.

It would be easier to just install another tool rather than improve those relationships, but the relationships determine whether the technology becomes part of how the business actually runs.

“Good enterprise AI should ultimately be visible in financial performance, operational metrics and customer experience,” says Xu. For her, model specifications, consumption measures and the number of tools deployed are insufficient proxies.

For the operator, useful measures include inventory turnover, fulfilment efficiency, costs and workforce productivity. For the customer, the question is whether needs are better understood and service improves. For partners, it is whether participation becomes easier and creates opportunities for growth.→



and partners to share capabilities and explore new AI applications across industrial supply chains.

Both show why applying AI in specialized environments requires more than technology alone. Domain knowledge, relevant data and real-world operating experience often sit across different organizations. Partnerships can bring these complementary capabilities together, helping AI move into real-world applications and enabling businesses to create value together.



AI often creates the greatest value not where there is the most information, but where information can ultimately be translated into decisions and actions."

→ Together, those measures connect internal productivity with external value.

But despite the potential of AI, Xu is clear that people are part of that calculation. JD.com trains frontline employees in robot maintenance and repair as automation spreads through its operations. Xu's view is that repetitive work can shift towards technology while employees develop technical skills and focus more on judgement and service – but that transition will require investment in capabilities, as well as equipment.

BUILDING OPEN ECOSYSTEMS

Xu expects the next stage of enterprise AI to depend increasingly on open ecosystems that bring together different types of expertise. JD.com's JoyInside embodied AI platform, powered by its proprietary JoyAI model, is one example. Through an open ecosystem, it is working with nearly 200 brands across appliances, robotics, healthcare, toys and other categories to bring AI into a wider range of hardware and real-world applications. JINGDONG Industrials' Baichuan Initiative takes a similar approach, bringing together industri-

al companies It's a lesson Xu believes other enterprise leaders can learn from: if you understand the business problem, connect the people and processes needed to solve it, you can then measure the result. AI becomes consequential when it changes the decisions that move products, fulfil needs and create value – but the real test starts when the demonstration ends and the operation is underway.


IMPLE- MENTING AI

For most enterprises, AI implementation has proven far more difficult than anticipated. Pilot projects proliferate, but scaling to production remains elusive. So how do you do it?


WHERE'S MOVING FASTEST ON AI?

Which areas are adopting AI in earnest, and what function sees the biggest gains?


DOES GENERATIVE AI ASSIST HIGH – OR LOW – EXPERTISE WORK?



GenAI Adoption Rates by Occ. (3-digit SOC)




GenAI Adoption Rates by Task (DWA)




TALENT STRATEGY ADJUSTMENTS DUE TO AI (%)


Question: How is your organization adjusting its talent strategies because of the adoption of AI tools/capabilities?



AI ADOPTION METHODS: 2024



AI ADOPTION METHODS: 2025




■ BUILT INTERNALLY ■ PURCHASED

(SOURCE: MENLO VC, THE STATE OF GENERATIVE AI IN THE ENTERPRISE, 2025)

ENTERPRISE AI MATURITY BY ORGANIZATION REVENUE, 2025

SOURCE: MCKINSEY & COMPANY SURVEY, 2025 | CHART: 2026 AI INDEX REPORT



GOVERNANCE, TRUST AND THE AI-READY ENTERPRISE

Overhauling an organization for agentic AI requires far more than tooling – it needs a vibe shift.



AS AI SYSTEMS BECOME MORE CAPABLE AND MORE

deeply embedded in enterprise operations, questions of governance and ethics start becoming much less theoretical and get down to brass tacks and practical necessities.

How do we ensure AI systems make fair decisions? How do we maintain accountability when algorithms act autonomously? How do we protect privacy and security in AI-powered organizations? How do we comply with rapidly evolving regulatory requirements?

Enterprises are grappling with them all today, as AI systems make hiring recommendations, determine creditworthiness, provide medical advice and execute financial transactions. For organizations, the stakes are high: poorly governed AI can lead to discrimination, regulatory penalties and tank the reputation of firms.

Overhauling an organization for agentic AI requires changing the internal tooling. But it also requires governance, data, talent and regulatory scaffolding that most organizations don't have today. →

→ AN EVOLVING LANDSCAPE

Enterprises operate well, but can be compelled to act further by regulatory incentives. And that regulatory environment for AI is taking shape around the world. From the European Union's AI Act, which establishes a risk-based framework for AI regulation, categorizing AI systems by their potential harm and imposing stricter requirements on high-risk applications, to executive orders on AI safety and security issued in the United States, there are a number of rules to follow. In the US, individual states are also passing their own AI regulations.

FROM COMPLIANCE TO ASSURANCE

Regulation sets the floor; enterprise governance has to go further. As AI systems gain access to data and tools, cybersecurity and model risk start to converge. Companies need to know what an agent can reach, what it is allowed to do, how its behaviour is logged and how quickly it can be stopped when something goes wrong.

That pushes assurance upstream into procurement, as well as deployment. Buying an AI system means testing not just capability and price, but data handling, security, model provenance, auditability, fall-back arrangements and the controls a supplier can demonstrate. High-risk systems may need independent evaluation before they are trusted with sensitive decisions or actions.

Boards do not need to approve every model, but they do need clear accountability for AI risk: who owns it, what thresholds trigger escalation and what evidence shows controls are working. The organisations that scale AI safely will treat assurance as an operating discipline, not a compliance exercise bolted on after deployment.

WHAT GOOD GOVERNANCE LOOKS LIKE

Effective AI governance requires several components working together. At its core you need model transparency

and explainability. It's vital to be able to outline how AI systems make decisions, to explain them to regulators, customers and affected individuals. That matters as much for business accountability as it does for regulatory demands. Can you explain why the AI made a particular decision? Can you demonstrate that the decision was fair and reasonable?

Linked to that is the need for bias and fairness assessment. Regularly testing AI systems for discriminatory outcomes and implementing mitigation strategies when biases are identified is crucial. Human oversight also ensures humans remain in the loop for high-stakes decisions and can override or appeal AI decisions – particularly important for agentic systems that can act autonomously within defined parameters.

Data governance teams also need to leap into action to protect the privacy and security of training data and ensure compliance with data protection regulations. This includes data lineage tracking, privacy impact assessments and security controls for training and inference data. Underpinning that are accountability structures that establish clear lines of responsibility for AI systems, from development to deployment to monitoring. Who owns the AI system? Who is accountable when it fails? Who has the authority to shut it down?

That feeds into ongoing monitoring and auditing that continuously assesses AI system performance and impact, making adjustments as needed. This includes technical monitoring for performance drift, business monitoring for ROI, and compliance monitoring for regulatory adherence. Despite the best hopes of users, AI systems are not “set and forget” – they require continuous oversight to ensure they remain aligned with business objectives and regulatory requirements.

FINDING TALENT

You might think an AI-enabled enterprise doesn't need to consider the human element. But building →

→ internal AI capabilities is essential for successful AI deployment. This doesn't mean every enterprise needs to hire hundreds of AI researchers, though it does mean every enterprise needs a balanced set of AI capabilities.

AI translators to bridge the gap between technical teams and business stakeholders, translating business problems into AI solutions and AI capabilities into business value, and AI engineers to deploy, monitor and maintain AI systems in production, are both crucial to the smooth running of AI systems.

At the same time, AI governance specialists within an enterprise can help establish and maintain the frameworks that ensure AI is deployed responsibly, compliant with regulations and aligned with enterprise values. They understand the regulatory landscape, risk management frameworks and accountability structures needed to govern AI effectively.

But organizations don't just need AI experts. Domain experts with AI literacy dovetail with other workers to identify high-value use cases, evaluate AI vendors and oversee AI deployments within their functional areas. They don't need to be AI experts themselves, but they need enough AI literacy to ask the right questions and make informed decisions about AI adoption.

CHANGING CULTURE

Perhaps more importantly than who you're hiring and what tech you're running, AI-ready enterprises require cultural change up and down the organization. Every employee who interacts with AI systems must understand their responsibilities for ethical AI use. Tech teams need to build governance into AI systems from the start, not bolt it on as an afterthought. And senior leadership must make AI governance a strategic priority, not merely a compliance exercise.


Such a significant culture transformation doesn't happen overnight. It requires leadership commitment, investment in training and education, and a willingness to have difficult conversations about AI's impact on the workforce and society.

But it's vital, whatever happens, to ensure that governance is more than just an afterthought. It is the foundation on which trustworthy AI is built. The enterprises that understand this will be the ones that succeed in the AI era. ■




WHAT IS THE STAGE OF GENERATIVE AI ADOPTION IN YOUR ORGANIZATION?

ALL RESPONDENTS



DOES YOUR ORGANIZATION HAVE AN AI GOVERNANCE OFFICER, OR AI GOVERNANCE OFFICE?

ALL RESPONDENTS



(SOURCE: PACIFIC AI GOVERNANCE SURVEY, 2025)

59%

of organizations have a formal role or office for AI governance

36%

of small organizations have a formal role or office for AI governance

62%

of large organizations have a formal role or office for AI governance

64%

of medium-sized organizations have a formal role or office for AI governance





HOW CAN FII INSTITUTE HELP?

Convening ideas and individuals to provide solutions.





AT FII INSTITUTE, WE RECOGNIZE THAT THE AI revolution is more than technology – it’s about people, organizations and societies adapting to profound change. Our mission is to serve as a trusted partner and convener, helping our members and the broader community navigate the opportunities and challenges of enterprise AI.

We are not just a think tank – we’re a “do tank.” We don’t just produce ideas; we convene the people who can turn those ideas into action.

CONVENING POWER

One of our greatest assets is our ability to convene the great and good of the world at our regular meetings around the globe. Our events bring together heads of state, CEOs, investors, entrepreneurs and thought leaders for substantive conversations that lead to real outcomes.

Our annual flagship summit convenes in Saudi Arabia each October, bringing together thousands of participants from more than 90 countries. This is where global leaders come together to debate the most important issues of our time – including the future of enterprise AI.

Our FII PRIORITY summits in the United States, Europe and Asia focus on regional opportunities and challenges, enabling deeper dialogue on specific topics – such as AI adoption in different regional contexts.

DATA-DRIVEN INSIGHTS

And that’s important because our work is grounded in data. Citizens in the Global South are much more positive

about the role of AI in their lives. They are more optimistic that AI will create more jobs than it destroys (58% vs 35% in the Global North). However, they share the concerns of citizens of the Global North that the use of AI in education will increase the digital divide (76% of global citizens overall are concerned about this) and are as likely to think that AI tools are a threat to society (54% of global citizens agree). Given these concerns, it is understandable that citizens across the world want to see a global framework to regulate AI (87%).

That data comes from the FII PRIORITY Compass, our global survey of citizen sentiment, which surveys over 50,000 citizens across more than 24 countries. This data grounds our discussions in what people actually care about.

When it comes to AI, the Compass reveals important insights: growing demand for AI sovereignty, concerns about AI’s impact on work and optimism about AI’s potential to solve real problems. These insights inform our programming and help our members understand the broader context in which they’re making AI decisions.

JOIN THE MOVEMENT

We invite you to engage with our AI work. Join us at an upcoming event. Participate in our working groups. Explore our research. And let us know how you are progressing on your AI journey.

Together, we can ensure that the AI revolution serves humanity – and that the benefits are widely shared. ■



KEY TAKEAWAYS TO TRACK

How should enterprises think about and deploy AI now and into the future? Here are some suggestions.



FROM LEADING LIGHTS IN INDUSTRY TO WORLD

leaders, FII Institute convenes the great and good from around the world to discuss the needs, stresses and strains we all face.

And how much to use AI, or not, is a key question many leaders in private and public sectors are facing. So what are the key messages to take away from this report?

1. AUTONOMOUS FUNCTIONS MOVE CLOSER

Agents are shifting from answering questions to completing bounded workflows. The next step is not a fully autonomous company, but more business functions in which software can plan, act and escalate within clear limits.

2. THE WORKFORCE BECOMES HUMAN + AGENT

The unit of work changes when agents can handle research, coordination and routine execution. Human roles move towards judgement, exception handling, supervision and the work where trust, creativity or accountability still matters most.

3. INVESTORS WATCH ADOPTION QUALITY, NOT AI LABELS

The useful signal is not whether a company says it uses AI, but whether it can move systems into production, redesign workflows around them and show durable gains without letting integration, governance or risk costs overwhelm the benefit.

HAVE YOUR SAY

Got big thoughts about how to deploy AI across your enterprise? We're always open to conversations and real-world ramifications. Join the FII Institute community at <https://fii-institute.org>.

There, you'll find information about all our latest events around the globe and our most recent white papers and Impact reports, plus recaps of our in-person events.



BUILDERS MOVE UP THE STACK

Models are only part of the opportunity. The next generation of enterprise AI companies will compete on orchestration, integration, life cycle management, assurance and the coordination layers that make agents useful inside real organizations.

4.

INFRASTRUCTURE BECOMES STRATEGY

Compute, cloud, inference, data and deployment choices set the cost, control and resilience of enterprise AI. Businesses will mix public, private, sovereign, proprietary and open-source components according to the workload rather than choose one stack for everything.

5.

ADVANTAGE COMES FROM EXECUTION

Access to capable AI is spreading quickly. The harder advantage to copy is organizational: knowing what to automate, what to keep human, what to build or buy, and how to change incentives, processes and governance so AI compounds rather than creates another layer of tools.

6.

AN ENTERPRISING IDEA

 **ENTERPRISE AI IS RAPIDLY** moving from experiment to everyday infrastructure. The technology is becoming cheaper, more reliable and easier to plug into existing systems. It all means companies are no longer just asking whether AI can write emails or summarize documents, but are integrating it into systems to search knowledge bases, support customers, analyze contracts, spot patterns in data and automate routine work that slows people down.

At its best, enterprise AI can give staff back time, improve decision-making and make services more responsive. It can help teams find information faster, reduce repetitive

admin and reveal insights that would otherwise stay buried in spreadsheets, inboxes or legacy systems. Used well, it becomes a productivity layer across the organization, not just a shiny tool for early adopters.

But mainstream adoption also raises the stakes. Businesses need to know where data goes, how models reach conclusions, who is accountable when they go wrong and how bias or security risks are managed. The opportunity is real, but only if adoption is deliberate, governed – and human-led.

Ensuring humans aren't just in the loop, but are considered, consulted and have ownership over how AI adoption in the enterprise takes place is vital

THE FII INSTITUTE IS GUIDED IN ALL IT DOES BY A STRONG PURPOSE, VISION AND MISSION:



PURPOSE

Enabling a brighter future for humanity



VISION

Empowering the world's brightest minds to shape a brighter future for ALL, and with ALL



MISSION

Curating and enabling ideas to impact humanity sustainably

FII-I has three pillars to deliver its mission:
THINK, XCHANGE and ACT

1 FII-I THINK

Identify societal challenges and current inhibitors

Curate the brightest ideas to address societal issues

3 FII-I XCHANGE

Create platforms for live discussions on the future of humanity

Share knowledge, stories and publications with different stakeholders

2 FII-I ACT

Catalyze innovation and initiatives by mobilizing partners and resources

– not for their own protection, but to ensure that organizations use AI well and for the benefit of their operations, rather than running counter to it.

Those that do this well will see AI embedded in workflows, tools and decisions people already use every day. The winners won't be the organizations that rush to automate everything overnight, but those that identify where AI genuinely improves work, train staff properly and build safeguards for proper use from the start. That means treating AI as a capability to manage, not a magic box to trust blindly on its own. The history of AI has shown us that claimed capabilities and actual reality can sometimes differ.

BOARD OF TRUSTEES



H.E. YASIR O. AL-RUMAYYAN
Chairman of the Board of Trustees, FII Institute; Governor, PIF



H.R.H. AMBASSADOR REEMA BANDAR AL-SAUD
Ambassador of the Kingdom of Saudi Arabia to the United States of America



H.E. MOHAMED ALABBAR
Founder & Managing Director of Emaar Properties



PROFESSOR ADAH ALMUTAIRI
Professor Pharmaceutical Chemistry, University of California (UCSD)



MR. RICHARD ATTIAS
Chairman of the Executive Committee, FII Institute; Founder & Chairman, Richard Attias & Associates



PROFESSOR TONY CHAN
Former President of King Abdullah University of Science and Technology (KAUST)



DR. PETER H. DIAMANDIS
Founder X-Prize Foundation & Singularity University



SIR NOEL QUINN
Chairman of Julius Baer Group



SENATOR MATTEO RENZI
Senator of Florence; Former Prime Minister of Italy

SOURCES

Overview of key sources used for this report

Alice Labs, AI Adoption Index, 2026

BCG, How Agentic AI Is Transforming Enterprise Platforms, 2025

Deloitte, The State of Art in the Enterprise, 2026

Digital Applied AI Agent Productivity Statistics, 2026

Dynamo Software survey, 2025

FII PRIORITY Compass 2025

Gartner, Trending Questions on AI and Emerging Technologies, 2026

Global SWF Report 2026

IMF, AI Will Transform the Global Economy. Let's Make Sure It Benefits Humanity, 2024

McKinsey, The State of AI, 2025

Menlo VC, The State of Generative AI in the Enterprise, 2025

NBER, What work does generative AI do?, 2026

Pacific AI, AI Governance Survey, 2025

Stanford AI Index, 2026

IMPRINT

FUTURE INVESTMENT INITIATIVE INSTITUTE

PUBLISHER
Future Investment Initiative Institute
3884 Alakheel District Unit 9
Riyadh 12382 - 6613
Kingdom of Saudi Arabia

CONTACT
For clarifications and inquiries, kindly email:
info@fii-institute.org

EDITOR
Chris Stokel-Walker

LAYOUT & DESIGN
Janosch Wojcik
Carina Dylus



For details on how to become
a member, please visit
fii-institute.org/membership/

POWERED BY



FOUNDING PARTNER



VISION PARTNERS



STRATEGIC PARTNERS

