AI INSIGHT SERIES · ARTICLE 1

AI: Hype or Business Imperative?

Individual employees swear by the tools. Company balance sheets, for the most part, do not show it yet. What separates the organisations closing that gap from those still waiting for proof?

28 October 2027 · 15 min read

THE GAP AT A GLANCE

8 in 10
employees say AI has made them more productive
McKinsey, State of AI 2026

37%
say AI has contributed positively to EBIT
McKinsey, State of AI 2026

95%
of generative AI pilots delivered no measurable P&L impact
MIT NANDA, 2025

8%
have scaled multiple strategic AI initiatives enterprise-wide
Accenture

Walk into any boardroom in 2026 and someone will mention artificial intelligence within the first ten minutes. Budgets have shifted, job titles are changing, and barely any strategy document goes out without an AI section. But beneath the enthusiasm, from boardrooms in Johannesburg to Silicon Valley, executives are grappling with a central question: Is AI merely the latest technology hype cycle, or has it become a genuine business imperative?

The evidence is mixed enough to support almost any argument you want to make. Individual employees swear by the tools. Company balance sheets, for the most part, do not show it yet. Understanding why that gap exists, and what separates the organisations closing it from those still waiting for proof, is the difference between AI as a genuine business imperative and AI as an expensive experiment in hope.

01

The Productivity Trap

McKinsey's 2026 State of AI survey found that eight in ten employees say the technology has made them more productive, and half say it helps them make better decisions. That is not a small number, and it is not a fluke. Across almost every level of an organisation, from the C-suite to individual contributors, people report genuinely different day-to-day experiences at work because of AI.

Yet ask the same organisations whether that individual benefit has shown up in the numbers that matter to shareholders, and the answer changes sharply. Only 37 percent of respondents say AI has contributed positively to their organisation's earnings before interest and tax, a figure barely changed from the year before despite a sharp rise in adoption.

MIT's NANDA initiative studied 300 enterprise AI deployments and interviewed more than 150 leaders in 2025 and found that 95 percent of generative AI pilots delivered no measurable impact on profit and loss, with only 5 percent translating into real, integrated value.

Boston Consulting Group's 2026 AI Radar, based on more than 2,300 executives including 640 chief executives, arrived at much the same conclusion from a different angle: 60 percent of companies report no material gains from their AI investment, and only 5 percent are generating value at scale. Three research organisations, three different survey designs, yet the same uncomfortable number.

This is not just a foreign phenomenon. Writing in IT-Online in April 2026, Accenture South Africa country managing director Kgomotso Lebele opened with a story about a major local bank running an AI-powered fraud detection pilot that delivered everything hoped for: faster detection, fewer false positives and measurable savings. Eighteen months later, the system was still running in the same contained environment it was built for. The pilot never became a platform.

> *Results were impressive: faster detection, fewer false positives, measurable savings. Eighteen months later, the system still operates in the same contained environment it was built for. The pilot never became a platform.*

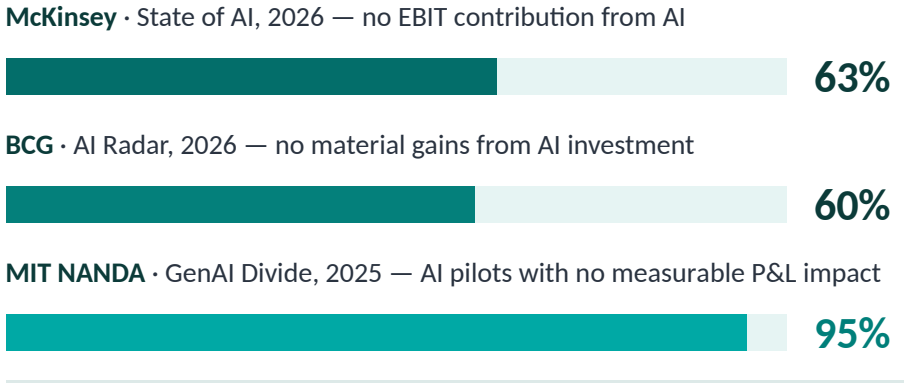
Kgomotso Lebele, Country Managing Director, Accenture South Africa

This is not just one bank's story. It is the shape of enterprise AI across South Africa right now, in financial services, retail, mining and telecommunications alike, and Accenture's research puts a number on how common it is: more than 63% of organisations plan to increase AI investment, yet only 8% have successfully scaled multiple strategic AI initiatives enterprise-wide. The other 92 percent are, in Lebele's words, running the same experiment on a loop.

This is not a new pattern for transformative technology. The lesson from the introduction of the computer age, and the one AI is now relearning, is that the first wave of any general purpose technology tends to deliver efficiency without necessarily delivering advantage. Chatbots, copilots and document summarisers help people finish tasks faster, but if every competitor has access to the same tools, the gains cancel out. Productivity becomes table stakes rather than a source of competitive advantage.

FIGURE 1

Three independent studies, one convergent finding



% of organisations / pilots. Each bar reflects that source's own survey question and methodology; figures are not directly interchangeable but point the same direction.

02

Where the Real Value Sits

According to McKinsey Quarterly research on where AI creates value, the technology moves through three overlapping waves, and leaders who manage only the first wave will be disappointed with the results.

1

WAVE ONE

Productivity

Tools like JPMorgan's real time fraud scanning or BMW's automated quality inspections speed up existing workflows without changing them. Useful, but rarely a moat, because rivals can buy the same capability.

2

WAVE TWO

Differentiation

AI enables new products, services or business models that were not previously possible. Legal technology firm Harvey blends generative AI with proprietary legal workflows to deliver work that once required teams of lawyers. Education platform Khanmigo offers personalised, always on tutoring that challenges the economics of the traditional classroom. Amazon Prime shows how AI driven personalisation and logistics, layered onto a subscription model, can shift competition from single transactions to long term relationships. These are not incremental improvements. They change what a company sells and how it earns money.

3

WAVE THREE · MOST DISRUPTIVE

Lower transaction costs and genuine business transformation

Largely through agentic AI systems that can search, negotiate and act on a person's behalf with very little friction. Imagine a case where AI agents compare energy tariffs and switch a household's provider automatically, or reallocate idle bank deposits toward the best available interest rate.

McKinsey estimates that a third of the seventy trillion dollars sitting in global consumer deposits earns close to zero interest, and shifting even a small share of that toward better rates could cut banks' deposit profits by twenty percent or more. This is not a hypothetical risk. It is the kind of structural shift that decides which companies still exist in ten years.

The historical parallel is electricity. Factories did not become dramatically more productive when they simply swapped a steam engine for an electric motor in the same layout. The real transformation came later, once managers redesigned the factory floor entirely around distributed power. AI is following the same arc. The productivity gains are the electric motor bolted onto the old machine. The advantage comes from redesigning the machine.

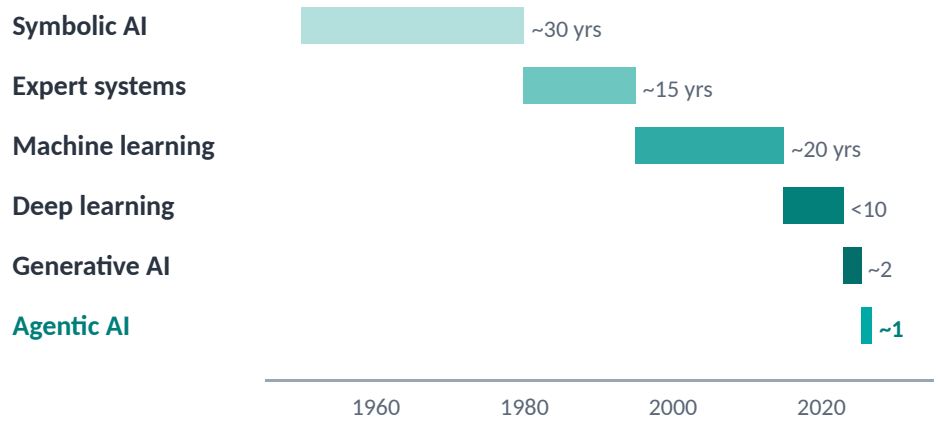
03

Why Waiting Is the Riskier Bet

It is worth giving consideration to how fast the ground has actually moved, because it changes the answer to whether an organisation can afford to adopt a “wait and see” attitude. AI's development runs through six distinct eras, and each one has arrived faster than the one before it. Symbolic AI, hand written logical rules with no ability to learn, held the field for roughly three decades from the 1950s. Expert systems, where specialists' knowledge was coded into elaborate rule books, took about fifteen years to run out of road. Machine learning, the point where systems began learning patterns from data rather than following fixed instructions, spanned two decades. Deep learning compressed that to under ten years. Generative AI, the technology behind the current wave of enterprise spending, took roughly two years to move from research curiosity to boardroom priority. Agentic AI, systems that take multi step action with some autonomy, arrived within about a year of that and is already being piloted inside large enterprises.

FIGURE 2

Each wave of AI has arrived faster than the last



Approximate years each era took to mature before the next layered on top. Durations synthesised from standard AI-development staging (symbolic → expert systems → machine learning → deep learning → generative → agentic).

Independent research confirms just how unusual that speed is. Stanford's 2026 AI Index found that generative AI reached 53 percent population adoption within three years of ChatGPT's launch, a faster climb than either the personal computer or the internet managed at the same stage of their

own lives. The same report found that AI agents' success rate on realistic, real world computer tasks jumped from 12 percent to 66 percent in about eighteen months, closing most of the gap with human performance on the same tasks.

Each stage did not really replace the one before it so much as get layered on top of it, which is why most enterprise AI today is a stack: pattern recognition underneath, generative output in the middle, and increasingly an agentic layer deciding what to do with it. The practical implication is not that the technology will slow down long enough for latecomers to catch their breath. It is the opposite. The gap between an AI category being emerging and being mainstream keeps shrinking, and the survey data backs that up: the share of organisations scaling AI across the enterprise rose from 38 percent to 44 percent in a single year, and large enterprises are now moving into agentic AI roughly twice as fast as smaller ones.

None of this means every agentic bet pays off, and it would be dishonest to pretend otherwise. Gartner expects more than 40 percent of agentic AI projects underway today to be cancelled by the end of 2027, mainly because of unclear business value and weak governance rather than any flaw in the underlying technology. The same research house found that roughly a third of organisations were still taking a formal wait and see approach as of early 2025. Put those two findings side by side and the lesson sharpens rather than softens: doing nothing is not a neutral choice, and doing something without discipline fails just as often. The risk was never speed itself. It is moving carelessly, or not moving at all.

None of that is an argument for chasing every new model release either. The organisations actually pulling ahead are rarely the ones with the newest tool. They are the ones that used the available time to build foundations that will outlast any single generation of technology: a small number of business domains chosen deliberately, a data foundation that can support whichever model comes next, and people who know how to redesign a workflow rather than simply operate a new interface. That is the honest answer to should we wait: no, but what an organisation does now should be building the capacity to move fast later, not betting the business on today's specific tool.

04

The Discipline That Separates Leaders From Followers

If enthusiasm alone determined outcomes, every company using AI would already be reaping the benefits. It does not, which is why the small group of organisations that actually shows a return, roughly one in twenty across three independent studies, is worth closer scrutiny.

McKinsey's version of this group, which it calls AI high performers, is more than three times as likely as its peers to plan for AI-driven fundamental business transformation over the next three years. Almost 75% of them report having fundamentally redesigned workflows because of AI, compared with barely 25% of everyone else. They prioritise growth and innovation goals alongside efficiency, rather than treating AI purely as a cost cutting exercise, and they invest more too, with over half planning to increase AI spending by ten percent or more in the coming year.

3× the cost reduction	**1.6×** the EBIT margin
~3× the return on invested capital	**FUTURE-BUILT COMPANIES VS PEERS** BCG AI Radar, 2026

Boston Consulting Group's independent research lands on a strikingly similar number by a completely different route. Its 2026 AI Radar found that so called future built companies, those that redesign operations around AI rather than layering it on top of existing processes, achieve 3x the cost reduction, 1.6 times the EBIT margin and nearly 3x the return on invested capital of their peers, while putting the share of companies generating AI value at scale at just 5%.

The common thread across all these studies is coherence rather than budget size. It is not that high performers spend the most. It is that strategy, workflow design, leadership commitment and risk management reinforce one another, rather than sitting in separate silos, each waiting for the others to catch up.

05

Proof From the Field

AWS had accumulated more than two hundred internal tools and forty disconnected systems over two decades of rapid growth, the kind of fragmented foundation many large organisations will recognise uncomfortably well. Rather than layering more AI agents on top of that mess, AWS spent time consolidating its data into a shared foundation before scaling automation. Once that foundation existed, an AI assisted pipeline analysis agent surfaced close to seventy-seven million dollars in additional commercial opportunities within a single business segment. Agent enabled workflows cut the time needed to prepare a customer strategy by around seven hours per engagement. More broadly, McKinsey research finds that organisations which pick one to three business domains and reinvent them end to end, rather than scattering AI thinly across the business, average a twenty percent uplift in EBITDA.

US$77m

in additional commercial opportunities surfaced by a single AI pipeline-analysis agent within one AWS business segment, after the company unified its data foundation.

Source: McKinsey, “Lessons on AI from our alliances,” 2026

06

An African Perspective

The sectors where South African companies have the most to gain from moving past the experiment loop are specific and already visible: real time fraud detection and payments automation in banking, a shift already made by 29 percent of banking institutions globally; AI-driven claims intake in insurance; demand forecasting in a margin-sensitive retail sector; predictive maintenance and safety systems in mining, where the payoff is measured in reduced downtime and protected lives, not just cost saved.

BANKING

Real time fraud detection and payments automation

INSURANCE

AI-driven claims intake

RETAIL

Demand forecasting in a margin-sensitive sector

MINING

Predictive maintenance and safety systems

PwC's Africa-specific research, published in May 2026 and covering 85 African companies within a global study of 1,217, finds the same pattern from a different angle. Africa's AI Fitness Index sits almost exactly at the global median, which sounds respectable until the detail underneath it is examined: the region trails AI leaders on every major dimension of AI-driven performance, and the gap is concentrated in execution rather than ambition or appetite. Twenty three percent of African CEOs who invested in AI reported revenue increases over the past year, and 25 percent reported cost reductions, real but modest numbers next to PwC's other finding, that the most AI-fit companies in its global sample generate 7.2 times greater AI-driven performance than everyone else.

None of that means the imperative case is theoretical in South Africa. Absa's Salesforce-built assistant, Abby, now operates in all 11 official languages and is used by roughly 15,000 Absa staff. Nedbank's own assistant, Enbi, is credited with cutting live agent chat volumes by more than half within a year of launch. Neither is a pilot sitting in a sandbox. Both are production systems that redesigned a real workflow rather than bolting a chatbot onto the old one. This is the key distinction, and it is playing out in Johannesburg just as it is in Seattle or Shenzhen.

07

Reading the Signals

Board level conversations about AI often get stuck on whether to buy a particular tool. AI behaves like a general-purpose technology, similar to electrification or the internet a generation ago, which means it will not show up as a single line item. It will touch cost structure, product design, the workforce and competitive dynamics all at once. The more useful question for a leadership team is not “should we buy this AI tool” but “which parts of our value chain are about to be restructured, and on what timeline”.

Three curves are worth tracking to answer that. The cost curve for AI compute and inference is falling steadily. The capability curve is rising, but unevenly across tasks and industries, so a model that is unreliable for legal judgment may already be excellent at drafting a first version of a contract. The adoption curve is accelerating across competitors, which makes timing a live variable rather than a background fact. Where an industry, and a specific company within it, sits on these three curves says more about whether **AI is hype or imperative** for that business than any single product announcement does.



Cost curve

Compute and inference costs falling steadily



Capability curve

Rising, but unevenly across tasks and industries



Adoption curve

Accelerating across competitors, so timing is live

That framing also clarifies the build, buy or integrate decision that trips up so many strategies. Building foundation models from scratch is rare and only makes sense for firms with unique data, scale and machine learning talent to justify the cost. Buying off the shelf tools for commodity efficiency is now the most common choice, and the right one for most productivity use cases. The highest leverage option, and the hardest, is integration: redesigning workflows around AI native tools in the specific places where proprietary data creates an advantage nobody else can copy. Defaulting to buy when the real opportunity is to integrate is one of the quieter reasons ambitious AI strategies stall without anyone noticing why.

08

Turning Intent Into Outcomes

Part of the reason the gap between individual productivity and enterprise impact persists is that most organisations still treat AI as a collection of pilots rather than a strategy. The symptoms are easy to spot: successful proof of concept projects that never make it to production, several departments independently buying similar tools, data scientists hired without the infrastructure to do useful work. Left unmanaged, these patterns waste a large share of AI spending on projects that quietly stall before they ever reach production.

A second trap is believing every capability needs to be built in house. Most organisations do not need to own the underlying technology that makes AI work. The businesses that pull ahead typically buy foundational infrastructure such as cloud platforms, security and orchestration tools from established providers, then concentrate their own effort on what genuinely differentiates them: proprietary workflows, institutional knowledge, governance and customer relationships. A real estate finance business that automated half of its invoice matching and billing workflow did exactly this. It bought the underlying agent development platform from a cloud provider, built its own workflow logic around it, and brought in an external partner for change management. None of that required reinventing the technology itself.

09

The Choice in Front of Leaders

The debate over whether AI is hype or a genuine business imperative is, in a sense, the wrong debate. Both are true at once. There is real hype, visible in tools that improve individual convenience without changing a single line on the income statement. Despite the hype, AI has crossed a threshold where ignoring it is no longer a viable strategy. There is a genuine imperative, visible in the handful of companies quietly redesigning how their businesses work and earning outsized returns for it. Organisations with rich, well-governed data assets can build competitive moats that are difficult for rivals to replicate. AI transforms data from a passive asset into an active driver of value creation. AI automates repetitive tasks, accelerates decision-making, and enhances human productivity. Many research bodies estimate that AI will add trillions of dollars to global GDP by 2030. These gains are not theoretical; they are already visible in industries such as finance, retail, logistics, and professional services.

The task facing leaders is not to decide whether AI matters. It is to decide which of these two categories their own organisation currently falls into, honestly, and to close that gap on purpose rather than hoping momentum will do it for them. That starts with a simple question worth raising at the next leadership meeting: are we using AI to work a little faster, or are we willing to redesign how the work gets done at all? The answer determines whether AI becomes this decade's electric motor, bolted onto the same old machine, or the moment the factory finally gets rebuilt.

THREE QUESTIONS FOR YOUR NEXT LEADERSHIP MEETING

1. Where does our AI activity actually sit: productivity, differentiation, or transaction-cost reduction and business transformation?
2. Which one to three domains would we redesign end to end if we were serious about it?
3. Are we buying tools, or building the proprietary advantage nobody else can copy?

REFERENCES

1. IDC, *The Global Impact of Artificial Intelligence on the Economy and Jobs*
2. McKinsey / QuantumBlack, "The State of AI in 2026: On the Road to ROI" (August 2026)
3. McKinsey Quarterly, "Where AI will create value—and where it won't" (April 2026), Antoine Montard, Dago Diedrich, Tanguy Catlin
4. McKinsey, "Lessons on AI from our alliances: What AWS's agentic journey can teach CEOs about rewiring for AI" (August 2026)
5. McKinsey, "On AI transformations: The AI dilemma — build, buy, or partner?" (Ankit Kapoor)
6. MIT NANDA (Project NANDA), "The GenAI Divide: State of AI in Business 2025" (July 2025), Aditya Challapally, Chris Pease, Ramesh Raskar, Pradyumna Chari
7. Boston Consulting Group, "AI Radar 2026: As AI Investments Surge, CEOs Take the Lead" (January 2026; survey of 2,360 executives including 640 CEOs)
8. Gartner, press release (25 June 2025)
9. Stanford HAI, "The 2026 AI Index Report" (April 2026)
10. Kgomotso Lebele, Country Managing Director, Accenture South Africa, "From AI pilots to enterprise value: What it takes to scale," IT-Online, 23 April 2026
11. PwC, "Decoding ROI from AI in Africa" (PwC's AI Performance Study, Africa perspective), 15 May 2026
12. Absa "Abby" (Salesforce Agentforce) and Nedbank "Enbi", South African production deployments: https://iafrica.com/absa-becomes-first-african-bank-to-launch-agentic-ai-for-customers/; https://kasisto.com/press-releases/nedbank-launches-enbi-an-intelligent-digital-assistant-powered-by-kasistos-industry-leading-conversational-ai/

CONTINUE THE CONVERSATION

Is your organisation working a little faster, or redesigning the work?

Speak to Thamani's advisory team about where AI fits in your value chain.

Contact Us

thamani. ASSURANCE | ADVISORY | TAX

Your partner in Assurance, Advisory and Tax.

WHAT WE DO

Audit & Assurance
Technology Assurance
Internal Audit
Financial Advisory
Outsourced Finance & Payroll

OFFICES

158 Jan Smuts Avenue
Rosebank, Gauteng, 2122

32 Kloof Street
Gardens, Cape Town

CONTACT

Julius Baartman
Head of Internal Audit
julius@thamani.global
083 689 2915

Thuto Motsie
CEO
thuto@thamani.global
072 723 4878

© 2026 Thamani. All rights reserved.

Disclaimer | Privacy Notice | Cookie Notice | Terms & Conditions | PAIA Manual