kyndryl

AI Readiness Report

2025

Foreword

In the Kyndryl Readiness Report, a global survey of 3,200 business executives combined with exclusive data from Kyndryl Bridge, our AI-powered digital business platform, we explored an emerging paradox among business and technology leaders who express both confidence in their IT infrastructure and concern over its readiness to face future challenges. With artificial intelligence fast becoming a core component of the modern enterprise, we wanted to learn more about how this paradox converges with AI innovation— and how prepared leaders feel for an AI-powered future.

This additional report, featuring AI-specific insights from the Kyndryl Readiness Report, illuminates a stark divide between current confidence in AI and future readiness. While 86% of leaders reported confidence in their AI implementation and believe it is best-in-class, only 29% said that their AI is ready to manage future risks. According to another finding, most businesses are investing in AI, but only 42% reported seeing a positive return on their AI investments.

Common barriers to AI adoption include data privacy and security, uncertainty around traditional value metrics, and regulatory requirements. Organizations must also overcome modernization challenges and address talent deficits in AI and machine learning skills as they confront their AI readiness. Perhaps because of these challenges, leaders reported feeling least ready for AI implementation compared to other IT elements like cloud technology, underscoring the need for businesses to balance competing priorities as they address the foundational technology challenges that will position them for AI success.

In addition to these exclusive insights, we also share perspectives on reconciling the AI readiness paradox. As enterprises consider how AI can help them achieve their short- and long-term goals, they have an opportunity to build a foundation of strong operational security, establish enterprise-wide trust in AI systems, and place people at the center of AI design to unlock new value. At each step in their journeys, business and technology leaders can play a unique role in driving the change that will help AI live up to its transformative potential.

At a moment when applications of AI are reshaping industries around the world—helping banks more quickly detect fraud, optimizing production lines for manufacturers, predicting future healthcare outcomes, and more—businesses and governments need trusted partners who can help them turn technology advances into a competitive edge. As the world's largest IT infrastructure services provider, Kyndryl is at the forefront of AI innovation. We look forward to continuing our collaboration with customers and our broad ecosystem of partners to solve AI challenges and seize new opportunities—now and into the future.

Michael Bradshaw
Global Applications,
Data and AI Practice Leader at Kyndryl

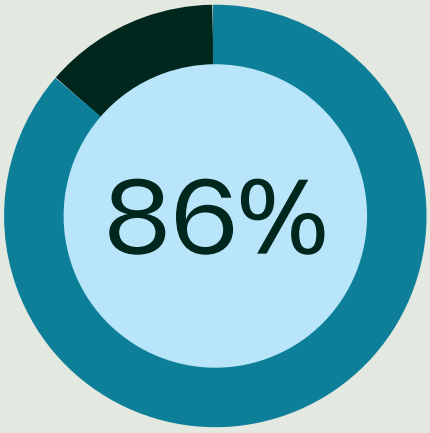

86% of leaders are confident their AI implementation is best-in-class.

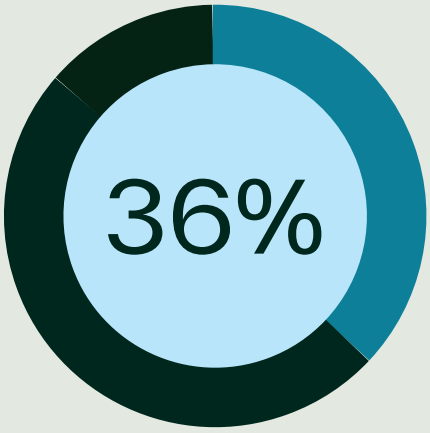

36% of leaders report ROI as a barrier to AI adoption

Three focus areas to improve security for AI projects in 2025

By Cory Musselman
Senior Vice President, Cybersecurity

Security professionals also can dream big about generative AI.

Wouldn't it be something if the technology could continuously analyze a company-wide set of telemetry—terabytes of data at a time from the myriad devices and identities and applications in the environment—and suggest when patterns move outside the normal standard deviation?

Imagine how that would revolutionize detection, protection and incident response capabilities—supercharging security teams' abilities to thwart bad actors and protect corporate and customer assets.

Yes, that's holy grail territory for a security leader.

A version of the capability already is operational. Security analysts can use generative AI to help collect, sort and do base analysis on incident data. It's just not yet at that grand, automated scale.

Such has been the story of what may be considered Wave 1 of this remarkable technology. Teams in every business function envision transformative use cases for generative AI. But most companies remain at the starting gate in terms of implementing it and realizing value.

Research for the Kyndryl Readiness Report found that concerns about security and data privacy are top barriers. The research also suggests only 29% of executives believe their AI capabilities are ready to manage future risks and disruptive innovations.

Most conversations about AI security right now focus on one of two themes. It's either how your organization can defend against bad actors who use AI to launch attacks and exploit data, or how security teams can use generative AI to improve the operational efficiency and efficacy of what they do.

Again, we're still in Wave 1. We must detach ideas about what generative AI may one day do from how we can responsibly use it today. It's where possible AI meets responsible AI. We have to host the two ideas at the same time, and work every day to close the gap.

To that end, I suggest three focus areas to improve security for AI projects in 2025.

1. Insist on AI governance

There's no overstating the importance of formalizing an AI governance framework. It's foundational work for any organization that intends to responsibly use generative AI.

Good governance not only helps ensure the right security, privacy and regulatory controls, but it also helps cut through the noise of ideas to determine how the company can get value from the technology in the near term.

At Kyndryl, we've established a cross-functional governance committee that evaluates and approves use cases, ensuring they align with regulatory requirements and the organization's risk tolerance. The group includes representatives from the security team, the CTO office, corporate strategy, legal and other functions.

In our case, through the governance process, a set of low-risk use cases has been prioritized, evaluated and approved. So if a team were to bring forward a proposed use that fits in one of those boxes, it's a quick approval. You know it's non-sensitive data. It's not regulated. It doesn't create unwanted exposure. At the same time, if a proposed use comes through but doesn't fit within the scope of what's already been approved, it has to go through the evaluation rigor.

So many organizations are trying to figure out how to use generative AI. Business teams feel pressure to adopt it, while security teams must ensure it's done safely and protects privacy. The formality of good governance isn't to squash innovation. It's to help balance speed to market and return on investment with ensuring it's done responsibly.

2. Reinvigorate cyber education programs

It used to be easier to spot a phishing attack.

Spelling mistakes were more prominent. The grammar just wasn't quite right. Unfortunately, that's a generative AI use case in the wrong hands.

Using generative AI, lower-level bad actors can now orchestrate more sophisticated phishing schemes and create more convincing deepfakes and other social engineering attacks. According to Gartner®, "by 2027, 17% of total cyberattacks/data leaks will involve generative AI."

To combat these kinds of attacks your company's first line of defense has to be employee training.

Securing and ensuring privacy is not just a security team or privacy office function. It's every employee's responsibility; particularly as human error is the No. 1 risk in protecting digital information. Herd immunity comes when each of us is cyber educated, aware and responsible for doing what we can to secure, protect and deliver results.

A challenge is that out-of-the-box cyber education programs often don't meet companies' specific needs. As threats evolve, so must cybersecurity education. At a minimum, it must become more customized, role-specific, engaging and timely so employees can relate the guidance to their daily activities.

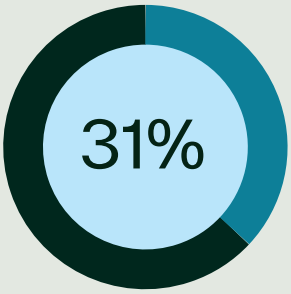

of leaders identified data privacy and security as a **top barrier to their AI adoption**

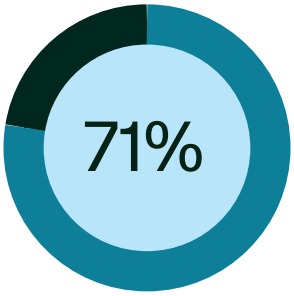

of leaders don't feel their **AI implementation is completely ready** to manage future risks

3. Consider a zero trust architecture

Also known as deny by default or never trust, always verify, zero trust models consider all traffic as untrusted. That's not a bad posture in an environment of increasingly more convincing deepfakes, for example.

Part of the strength of a zero trust model is in the many layers of security it encompasses—from the identity of the user and the device they use to the network, application and data they try to access.

The progressive validation process can derail nefarious activity when a first line of defense—such as voice recognition—has been compromised.

U.S. government agencies were already mandated to adopt a zero-trust architecture. McKinsey suggests that adoption rates for this architecture will continue to increase for companies of all sizes.

Zero trust is not a turnkey solution. It's a process, as much as an approach to fortify an organization's cybersecurity posture. But like formalizing governance and customizing training, moving to a zero trust model can help move the needle on readiness for the next wave of generative AI.

Gartner Press Release, Gartner Forecasts Global Information Security Spending to Grow 15% in 2025, August 28, 2024

AI's real challenge? Human behavior

Realizing the benefits of AI depends on human-centered design

By Victoria Pelletier
Vice President, Consult Partner

Business leaders everywhere are chasing the promise of artificial intelligence, captivated by its potential to unlock untapped benefits and opportunities. But the real test of AI's power won't be realized in bold headlines or grand visions—it'll unfold in the everyday routines of the workplace, as people reshape their behaviors around this transformative technology.

But realizing the true value of AI requires both scale and considerations beyond technology—and many organizations still lack the cultural, structural, and strategic readiness to make the best use of cutting-edge AI tools.

In our recent readiness report, we found that while 73% of business and technology leaders report investing in AI, only 41% say they are seeing a positive return on investment. It's a symptom of a thorny challenge. A quarter of leaders say they've run into difficulties integrating AI into their existing workflows. AI readiness at the enterprise level often falls prey to the allure of technological optimism, a belief that simply deploying advanced tools will yield immediate, transformative outcomes. Yet, as with many innovations, successful implementation and adoption depends not on the technology itself, but on how deeply it considers and integrates with the human systems and behaviors it aims to enhance. Integration is important, but we need to be clear that it's not just about accepting AI tools—it's about full-fledged user adoption.

At its core, AI readiness is a behavioral transformation—one that demands more than just new technology; it requires a strategic approach to change management. Enterprises must embrace human-centered design principles to navigate this shift successfully, ensuring AI solutions align with the needs, concerns, and expectations of those who will interact with them—employees, customers, and leaders alike. Effective change management plays a critical role in driving adoption, reducing resistance, and fostering a culture of trust and adaptability. It's about more than implementation; it's preparing people for the inevitable disruptions, addressing the exceptions to process or workflow standards, and providing the right support to help them embrace AI as an enabler rather than a disruptor.

It's crucial that leaders eager to adopt AI tools take the initial step of designing the desired experience with all the relevant personas in mind, which involves mapping out the nuanced needs of people—from end-user customers to frontline employees to tier 2 or escalation or exception management leaders. It's an exercise that requires thoughtfulness and an ability to think beyond the technology itself. Rolling out automated solutions into a workstream? It's easy to lean on technologists to build and deploy them. But what happens when an end user throws a curveball—an unexpected request, an unplanned detour? And when that inevitably happens, how do businesses avoid the added delays and frustrations that come with it? The answer lies in solid human-centered design—tackling these challenges upfront by bringing together a diverse group of stakeholders, not just the technologists, to create a more holistic solution that's built to handle the unexpected.

Consider as an example, a customer service chatbot. Its success depends not only on handling routine queries from customers, but also on ensuring that employees tasked with managing unexpected queries and escalations are empowered to resolve issues seamlessly. If the design overlooks the exception-management process—how and when problems are routed to human agents—the system creates friction instead of reducing it. Similarly, customers encountering repetitive loops in chat interfaces are likely to "zero out" to a human representative, undermining the very efficiency the technology was meant to deliver.

Secondly, achieving high-quality experience relies on alignment among senior leaders on a shared strategic vision, supported by clear communication and cross-functional collaboration. That type of readiness can be difficult to achieve in federated organizational models. Implementing AI in siloed environments is especially complicated, inviting instances of ill-defined commercial objectives, poor data quality, and inadequate integration with existing systems. Altogether, these create bad experiences.

Think about it through the lens of a memorable dining experience. Convincing customers to return to a restaurant isn't just about the menu, the same as getting repeat users to engage with a chatbot isn't just about functionality. It's also about service, the overall end-to-end experience. Customers and employees are more likely to engage with systems designed to feel and be intuitive, empathetic, and reliable. That starts with leaders having and executing on a shared strategic vision.

Successfully integrating AI into workstreams is challenging enough without business and technology leaders complicating the process due to poor communication and lack of commitment.

Finally, strategic alignment and persona mapping are critical for a third component of behavior change: establishing trust. That's an important detail, especially in organizations where some may be resistant to change. In companies where employees view AI with skepticism—or worse, fear—it is key that decision makers adopt "whole human leadership," which is all about linking technological change to the organization's broader purpose. It's also about demonstrating clearly how employees play a pivotal role in achieving stated goals.

Ultimately, AI readiness is not a purely technological challenge; it is a human one. By embedding human-centered design into every stage of implementation, enterprises can ensure their AI initiatives are not just operationally effective but also emotionally resonant, driving long-term adoption and meaningful outcomes.

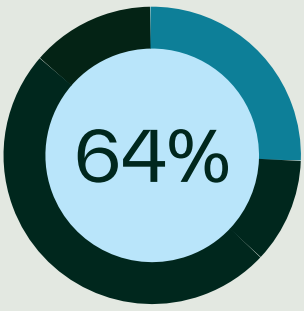

of leaders report feeling their IT is not prepared for AI implementation because of **a lack of IT skills and talent**

Customers and employees are more likely to engage with systems designed to feel and be **intuitive, empathetic, and reliable.** That starts with leaders having and executing on a shared strategic vision.

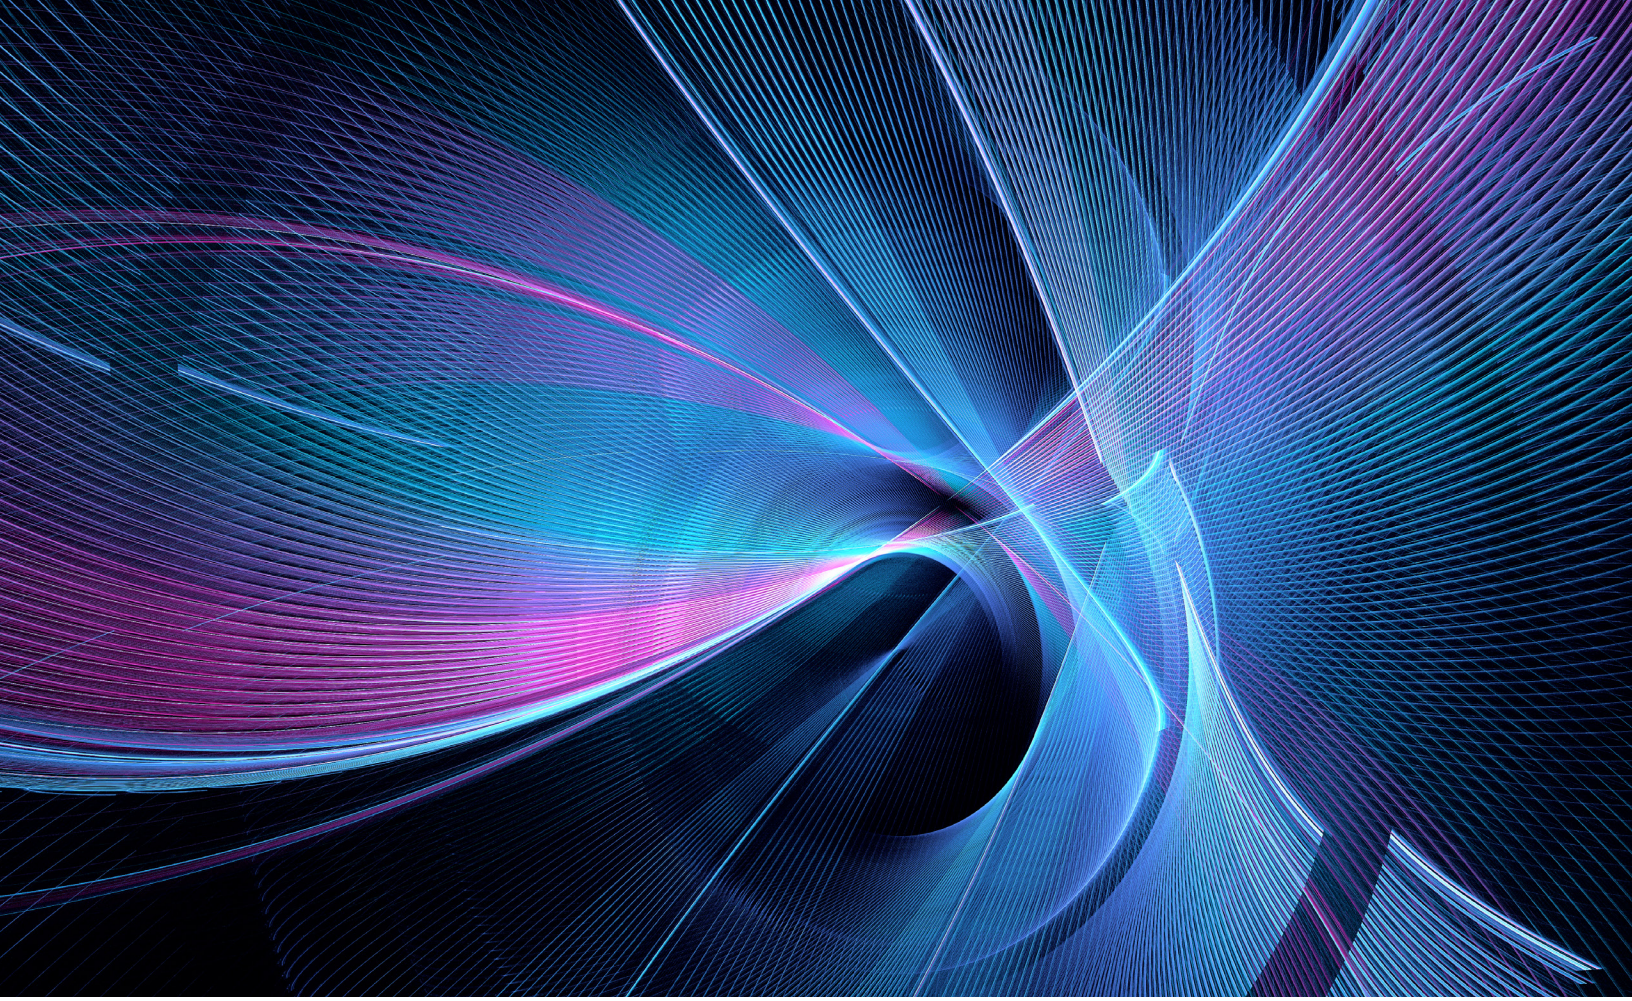

In the Wild West of AI, readiness begins with trust

By Kim Basile
Chief Information Officer

Artificial intelligence (AI) is reshaping industries at a relentless pace, and organizations must do more than just adopt new technologies; they must build trust to make them work. Trust is the linchpin of AI readiness. It's about transparency, communication, and empowering people to lean into change, not fear it.

When people hear the words "automation" and "efficiency," some often assume that their jobs are at stake. But the reality is that AI's true potential isn't in eliminating roles — it's in elevating them. Our focus has always been to free up teams from repetitive, low-value tasks so they can focus on work that requires creativity and strategic thinking.

Many leaders expect AI to fundamentally change their business in the near future, yet many still struggle with internal trust and adoption. A full quarter of enterprise leaders say they struggle with integrating AI into their operations, according to the recently-published Kyndryl Readiness Report. This underscores the importance of clear messaging and consistent reinforcement. When you help people understand AI as a tool to make their work more meaningful, you see the shift happen. It turns uncertainty into curiosity. But you can't just say it once and move on. The message must be reinforced in everything you do.

To that end, it's crucial to internalize the fact that trust doesn't happen by accident—it needs structure. That's why we've developed a governance framework that puts thoughtful boundaries around every AI initiative. Each proposed AI use case goes through a rigorous vetting process involving multiple departments. We want people to know that we're not throwing untested algorithms into the mix and hoping for the best.

Another noteworthy finding from the Kyndryl Readiness Report is that while 90% of business leaders feel confident about their current IT infrastructure, only 39% believe their systems are ready to handle future challenges. This gap reinforces why our governance approach is essential—it provides a roadmap for using AI responsibly, building trust along the way.

Lessons from Past Transformations

Through past transformative initiatives, I've learned that successful innovation requires a balance between ambition and practicality. It's not always about chasing the newest, flashiest tools but about refining existing systems to work smarter and more seamlessly.

We've used AI to integrate systems so they "talk" to each other, rather than operating as disjointed parts. This keeps operations streamlined and people focused. In today's rapidly evolving AI landscape, where there's always a temptation to start fresh with the latest tool, the real win is making your existing ecosystem smarter.

We've also found that the best AI use cases don't come from the top—they come from the people doing the work. That's why we empower employees to self-nominate tasks they believe could benefit from automation. When they see their ideas come to life, it's a powerful trust-builder. It tells them that we're not imposing change, we're collaborating on it.

Innovation happens when people feel like co-creators rather than bystanders. This approach has transformed how our teams engage with AI. Instead of asking, "What will this take away from me?" they ask, "How can this help me do my job better?"

25%

of leaders report difficulty integrating AI technologies **with existing systems and workflows**

The Power of Clear Communication

You can't build trust without communication. That doesn't mean delivering polished corporate memos—it means being open about the good, the bad, and the ambiguous. From the start of an AI initiative, we share what we're doing, why we're doing it, and what we hope to achieve. When things don't go as planned, we don't hide it. We explain it.

It's also important to connect the dots for people. Efficiency gains shouldn't feel like an end in themselves—they're a way to make space for more meaningful, high-impact work. When you show people how AI can lighten their load and make room for creative problem-solving, you turn skeptics into advocates.

No matter how advanced AI becomes, it's people who drive progress. Technology is only as good as the people behind it. We've made it a priority to embed this belief into our processes —whether it's through governance, feedback loops, or celebrating the wins that come from our people.

Leadership in the AI era requires more than technical expertise—it calls for empathy and transparency. If we tell our teams that AI will empower them, we need to show them what that looks like in real terms. For example, if we free up 25% of someone's time through automation, we help them use that time to grow, whether that means learning new skills or contributing to higher-level projects

Building the Future of AI Readiness

AI isn't just about having the latest tools. It's about creating a culture where change feels exciting, not threatening. Trust, transparency, and collaboration make that possible. We're not using AI to replace roles—we're using it to unlock new possibilities.

When people understand the "why" behind the technology, they're more than just users—they become advocates. By building trust at every step, we've created an environment where our teams don't just adapt to change—they help lead it.

We've also found that the best AI use cases don't come from the top—**they come from the people doing the work.**

kyndryl

Company Headquarters

One Vanderbilt Avenue, 15th Floor
New York, NY 10017

kyndryl.com

Kyndryl commissioned Coleman Parkes Research to survey 500 enterprises that rely on mainframes. This paper outlines the key findings of this survey and the implications for mainframe decision-makers.