Freeform Dynamics in association FUJITSU

Laying the foundations for enterprise AI

The importance of platforms and data

AI is currently the hottest topic in tech. But how far have enterprises progressed towards realizing the potential, and how should they now focus their efforts to ensure success? A survey of 400 data professionals provides some insights.

Key takeaways

- Enterprise AI expectations are high, with activity unfolding across the business, but most mainstream organizations indicate early-stage experimentation and adoption rather than broad established deployment.
- User-centric AI solutions like copilots and embedded features in business applications see strong acceptance, while autonomous AI agents face more skepticism, especially for complex decision-making.
- People-related challenges feature prominently when it comes to implementation concerns, with legacy mindsets and conventions cited as major impediments to strategic AI adoption.
- Data related issues represent the most critical barriers, however, with two thirds reporting data quality challenges and almost as many citing data availability problems as progress limiting factors.
- On-premises AI platforms are gaining traction alongside cloud services, driven by data sovereignty, a reluctance to move on-premises data unnecessarily, along with concerns over cloud costs and complexity.
- Fewer than a third of organizations have started to lay coherent infrastructure foundations to support their AI strategy, but those further along already show signs of being able to accelerate progress and time to value.

Why this research, and why now?

As we were planning this research in late 2024, the tech industry was heavily focused on the promise of AI, with a lot of emphasis on how AI models, specifically large language models (LLMs), were becoming more and more powerful. Debates were raging about who was winning a notional 'arms race' between the big tech players, and even between geopolitical continents on the world stage.

At the same time, excitement was growing at industry conferences around 'Agentic AI', and the possibility of building autonomous intelligent agents that could ultimately form the basis of a digital workforce to operate alongside human workers.

Meanwhile, our conversations with business and technology leaders working in a mainstream enterprise context were much more grounded. Many were struggling to translate all of the buzz and drama into practical, actionable initiatives that can deliver real business value.

Furthermore, the success stories we heard all seemed to revolve around AI solutions that leveraged enterprise data to provide relevant context and substance, and to avoid so-called 'hallucinations'. This in turn shone a spotlight on the platforms needed to deal with the sourcing, storage, movement and processing of data required for AI systems, along with real-world practicalities such as security, privacy, sovereignty, compliance, costs and the delivery of tangible benefits.

This conversation may not be as glamorous as media narratives about the AI wars and the threat to mankind, but it's far more important to any business looking to move forward with AI productively and confidently.

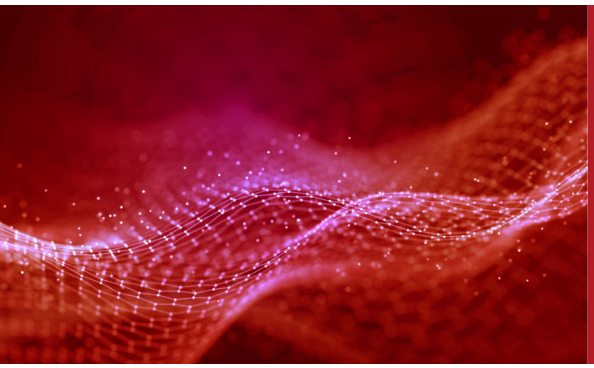

It can be hard to translate the AI buzz into practical initiatives that deliver real business value.

Research methodology in brief

The insights presented in this report are based on input gathered from 400 data professionals during interactive conversations over the telephone or via web conference. Interviewees were predominantly data scientists, business analysts, enterprise architects and senior IT/digital professionals (e.g. CDOs). Respondents were drawn from medium and large mainstream enterprises operating across a range of industry sectors. Geographies included DACH, UK, France, Spain, Nordics and Benelux.

Freeform Dynamics was responsible for survey design, data collection and analysis, together with results interpretation and the reporting of study findings.

The research was sponsored by Fujitsu.

AI has captured imaginations and expectations are riding high

It's no surprise to see our research confirming a significant level of activity and expectation around AI. Some organizations already report initial results in key functional areas, and the majority anticipate broader impact across the business, often within a short timeframe.

These findings reinforce the perceived broad applicability of AI and the notion that it will ultimately touch many aspects of operational management and business decision-making. A clear belief that it's a case of 'when' rather than 'if' comes through strongly.

Breadth of current and anticipated impact of AI across the business

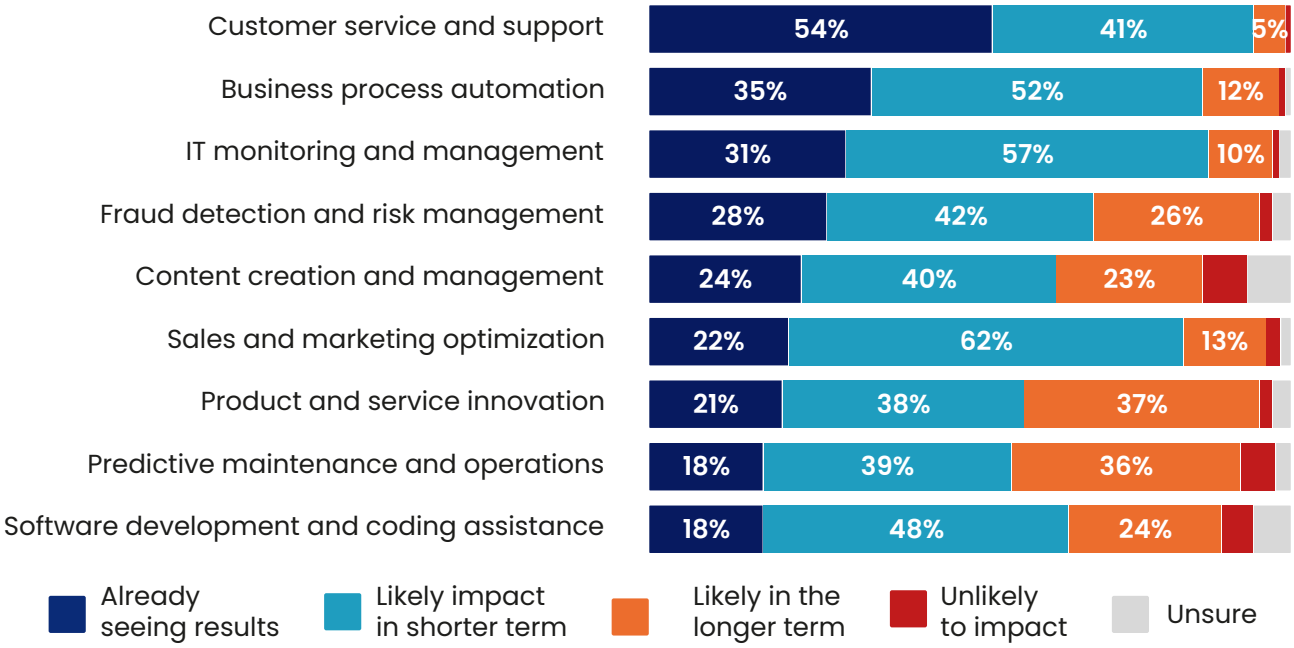

Before we get too carried away on a wave of positivity, however, we need to pause to consider what we mean by 'AI'. The views we see reflected in the above chart, for example, will have been formed from experience with more 'traditional' forms of AI as well as early experiences with LLM-based GenAI.

The truth is that way before ChatGPT and similar services burst into the limelight, various forms of machine learning had been used for quite a few years to enhance systems in all of the areas we see on the above chart. Indeed, most major software and SaaS providers have been embedding 'AI-powered' features into their solutions for well over a decade, marketing their AI credentials heavily as part of this.

With this in mind, much of the activity shown above will reflect the well-established use of machine learning in areas such as customer self-service, workflow automation, IT-Ops, security/risk management, and so on.

We might even put forward a counter-narrative that progress with AI is actually not that impressive given how long practical solutions have been available. A more useful way to look at it, though, is that interest in GenAI in particular has sparked or rekindled interest in other forms of AI, which is good.

Unless indicated otherwise, we use the term 'AI' inclusively throughout this document, i.e. to refer to all forms of traditional and emerging technology in this space.

Limited experience of success means many remain wary

With the exception of a relatively small number of skeptics, the use of AI to help employees do their jobs more effectively and efficiently is generally recognized and accepted. The areas we see listed here have become increasingly well-served in the early market, from copilots used in an office tooling context, through AI features embedded into CRM, ERP and other business applications, to AI enhanced analytics solutions. Many respondents will therefore have had some direct exposure to these solution types.

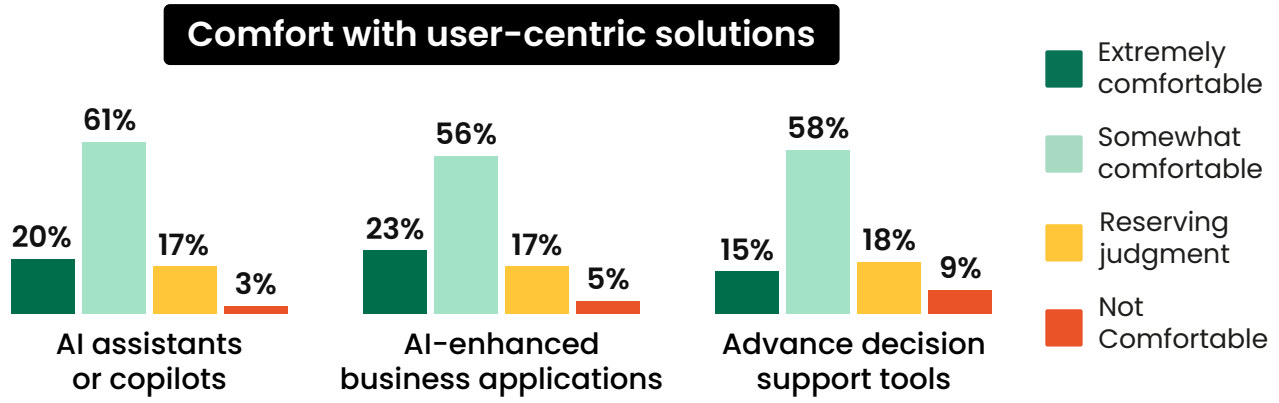

It's worth noting that where respondents have reservations (including those responding 'Somewhat comfortable'), this isn't always down to lack of trust in AI technology itself. As we'll see shortly, concerns and perceived challenges may be more to do with people and data.

Moving on to the hot topic of AI agents, the level of industry enthusiasm is not currently echoed among data professionals. While pure informational agents are reasonably well accepted, uncertainty rises when AI-enabled software or hardware robots are fully empowered to perform complex tasks.

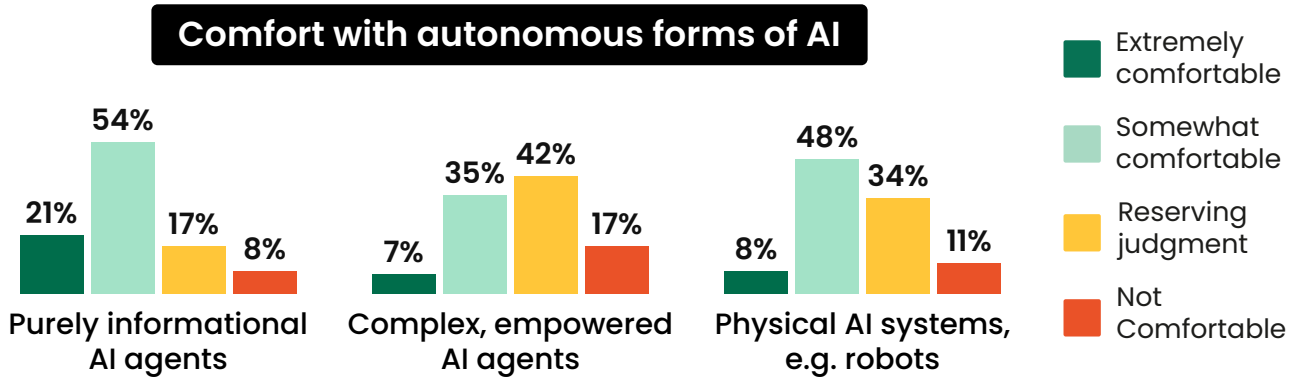

Reservations about empowered AI agents may stem from respondent's early hands-on experiences with chatbots, copilots and the like. The variability and unpredictability of responses observed with these translates to nervousness about how much AI can be trusted to work in a fully 'hands-off' manner. Whether this is fair depends on how 'Agentic AI' is defined and implemented, which remains the subject of much debate.

Progress depends on overcoming some fundamental challenges

Switching our attention to implementation practicalities, our research participants express a whole range of concerns.

Despite a feeling of limitless possibilities, there's still a shortage of well-defined solution patterns that can be taken off the shelf and simply implemented. It's therefore currently down to business managers, analysts and technical teams to collaborate to identify candidate use cases and assess the amount of tangible value they could deliver.

But even when you have identified a suitable use case, scoping the initiative to ensure deliverability and to manage stakeholder expectations amidst all of the industry promises represents the next set of issues.

A rigorous approach to distinguishing 'needs' from 'wants' is critical here. Only then can you combat the hype and work around the problem of people having been influenced by the seductive nature of AI tools encountered through personal use.

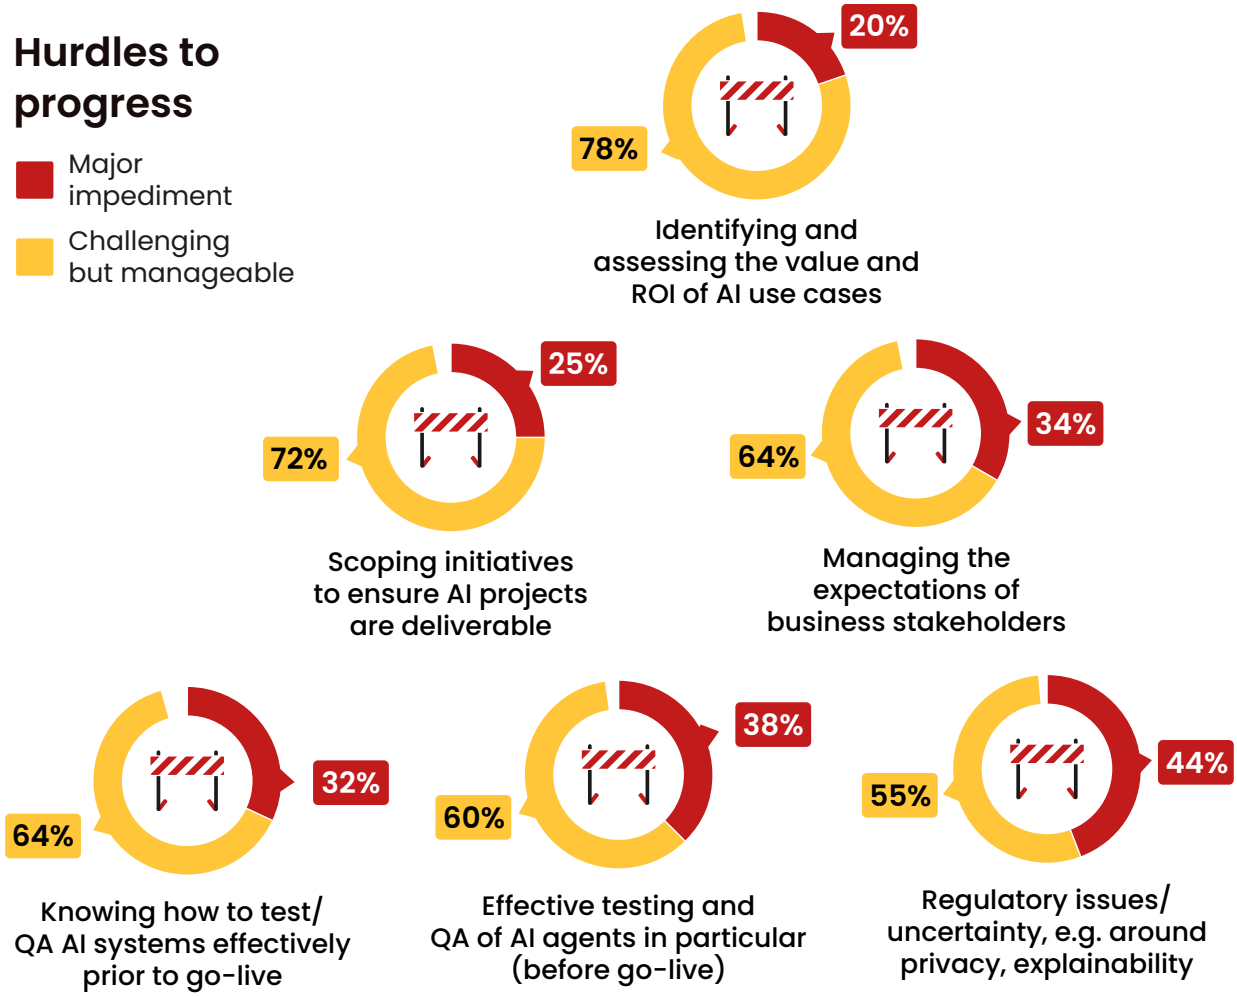

Building on earlier comments relating to LLM response variability, it's no surprise to see concerns around testing and QA appearing on the last row of this graphic. These sit alongside the topic of regulatory compliance and the uncertainties that frequently exist in this area. The rapid pace of change in the AI space creates additional challenges here as it can be hard to validate what's effectively a moving target.

Cultural and workforce issues stand out as critical to address

An obvious challenge we haven't yet mentioned is how to handle the potentially profound impact of AI on staff within the workforce. On the one hand, deploying AI powered solutions can totally transform the way some people work. On the other, getting the most from AI driven applications and using them safely can require mindset and behavior shifts. Some employees may find the change difficult or push back against it.

Hurdles to progress

- Major impediment
- Challenging but manageable

Overcoming legacy mindsets, conventions and habits

Employee resistance, e.g. trust issues, readiness to change

Following on from our earlier discussion of use case selection, ensuring good ROI is not the only reason for a rigorous approach. Another important consideration is the need to deploy AI in a purposeful rather than speculative manner. Anchoring initiatives in specific use cases is very useful here.

This principle is reinforced by the increasing number of stories emerging of copilot-style solutions struggling to meet expectations.

Experience shows that simply switching on an AI feature and leaving users to figure things out is often a recipe for disappointment, frustration and unnecessary cost and risk.

It's therefore not surprising to see the need for training and controls being widely recognized. Most organizations have already started on their workforce enablement journey, often building on previous efforts in the traditional AI space.

To what degree have you adopted the following?

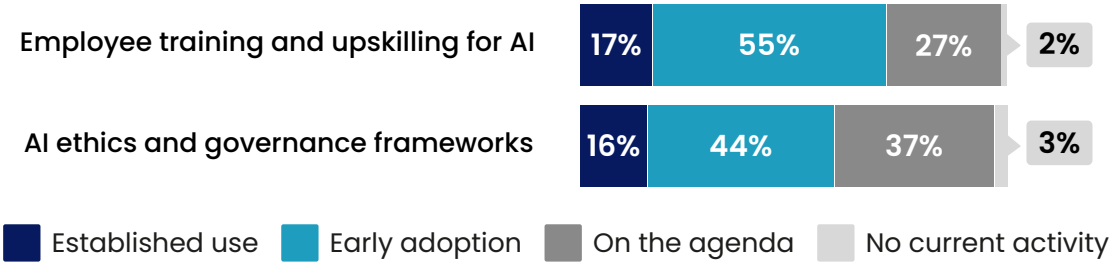

Whether it's boosting the level of AI literacy among knowledge workers, or retraining frontline operatives to work with AI-enhanced processes, appropriate guidance is clearly important, as is the need for training and controls on AI safety, ethics and governance.

Platforms are important, but cloud offerings will only get you so far

As with any other tech adoption, sustained success over the longer term is dependent on laying the right foundations as early as you can. In the case of AI, this means looking beyond ad hoc user-driven GenAI adoption or simply switching on AI features embedded in business applications. You have to consider the kinds of platforms you will need.

Such platforms deal with requirements ranging from data collation and preparation, through training and hosting AI models, to executing and orchestrating AI applications in a secure, scalable and resilient manner.

Looking at current activity, we see most established adoption revolving around cloud-based platforms, which is not surprising given the lower barrier to entry.

That said, it's notable to see the emergence of a respectable level of activity around on-premises platforms. Indeed if we look across the 'Established use' and 'Early adoption' responses on the chart below, the aggregate numbers are very similar. This suggests that private platform adoption is set to catch up with cloud through the use of on-premises or privately hosted infrastructure.

Use of AI platforms

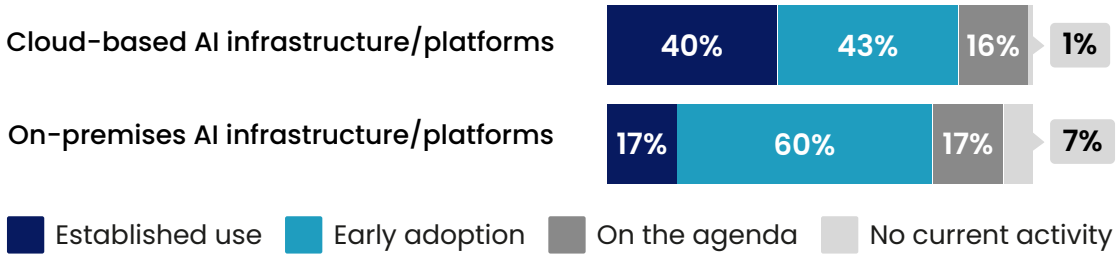

The picture we see here is consistent with the frequently reported 'cloud repatriation' trend, driven at least in part by falling hardware prices and the emergence of innovative payment models for on-premises infrastructure. There are then cost issues stemming from the volatility of the AI cloud infrastructure market, as shown below.

Platform/infrastructure related issues

But in addition to costs, there's another set of considerations at play that are arguably more important. These are associated with the data required to implement AI-based applications in an enterprise context, which as said earlier is a critical piece of the puzzle.

Enterprise data is key, especially structured on-premises sources

While the LLMs that underpin most current GenAI solutions are trained on huge data sets, your organization's own data is essential to provide context, as well as material to ground systems in your business reality. The research confirms this in no uncertain terms, with a range of different information types identified as important .

Data required to drive AI applications

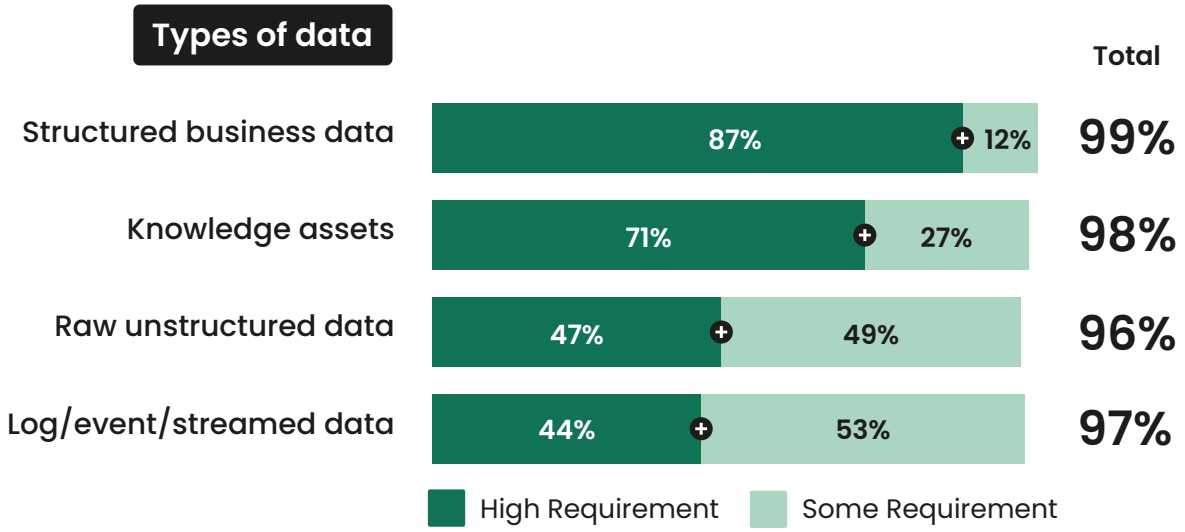

The most important type of input highlighted is structured business data. This makes absolute sense when you consider that customer transactions, operational data, financial records, etc provide a huge amount of historical and current context, as well as insights in their own right.

However, knowledge assets are not far behind, such as all of the information stored in document repositories and collaboration systems. Raw unstructured data such as various forms of messages, logs and streams that continuously flow around all businesses are also called out by research respondents.

Particularly interesting, and pertinent to the platform and infrastructure discussion, is the location of data. Here we see the emphasis on on-premises sources, which especially matters when considering sovereignty, performance and where processing occurs.

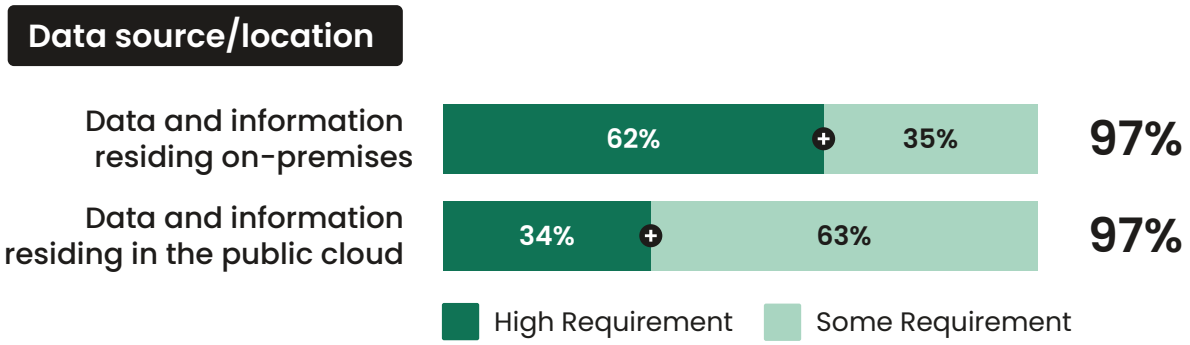

Data readiness for AI cannot be assumed; most have work to do

While it's easy to understand the need for enterprise data to enable success with AI, sourcing and preparing the data required is not always that easy. Data quality and consistency are common challenges, and together these represent a major impediment to progress for two-thirds of organizations. Related to this, simply sourcing and organizing data to feed into the AI environment is also a common challenge.

Progress limiting factors

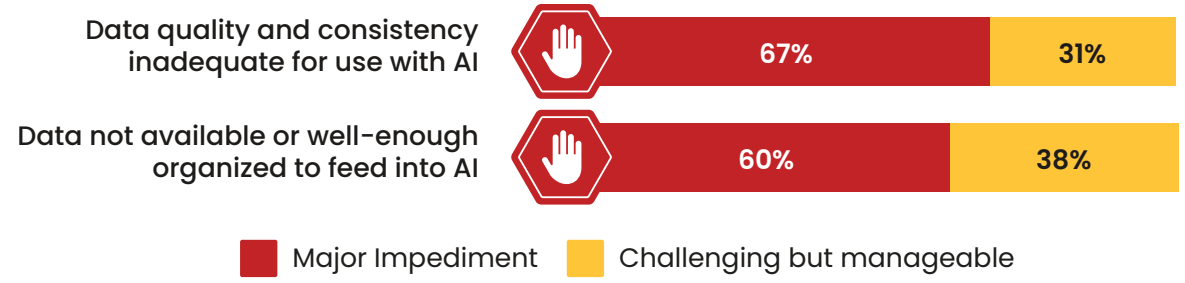

Drilling into some of the underlying causes and aggravators, we see a range of familiar challenges, many of which stem from fragmented data and ownership. It's also interesting to see the degree of reliance on legacy infrastructure called out by many. This suggests that modernization of information systems is an integral part of the AI discussion, especially when considering how to scale up the use of AI over time.

Causes and aggravators to navigate past

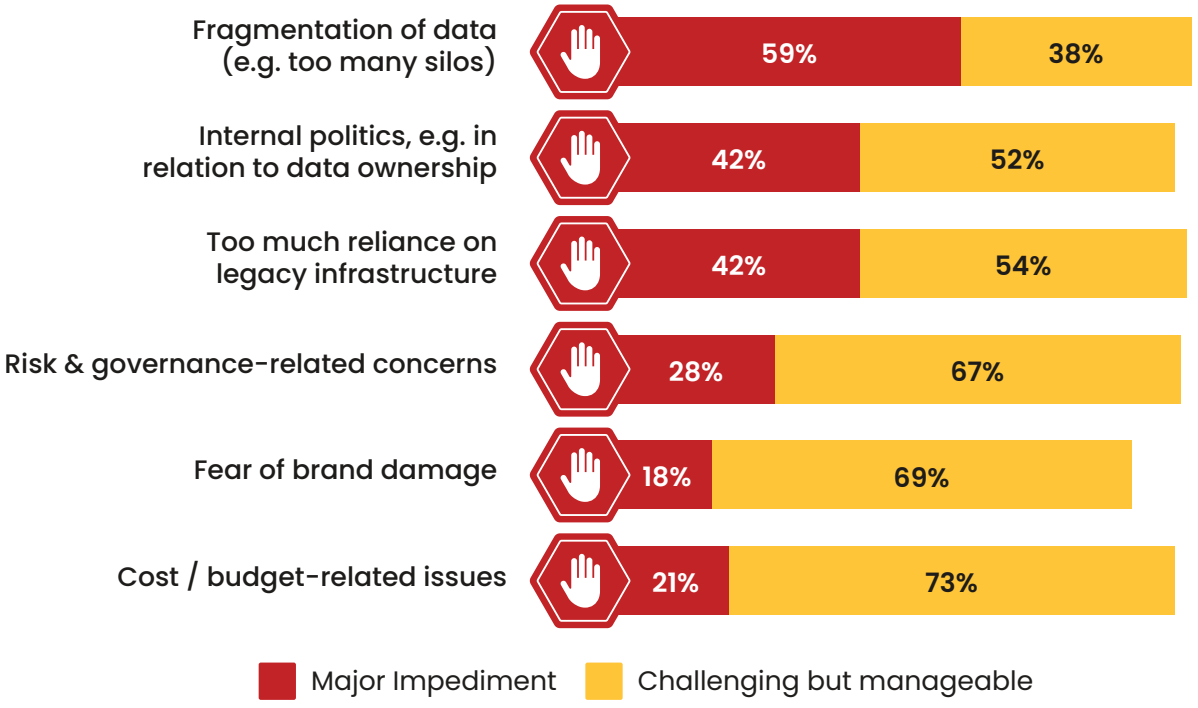

Laying strong foundations is critical for scaling your use of AI

We've already touched on the need for the right platforms in general, but picking up on some of the data related considerations, we see clear acknowledgment of the importance of some key capabilities to enhance AI with enterprise data.

Foundational capabilities

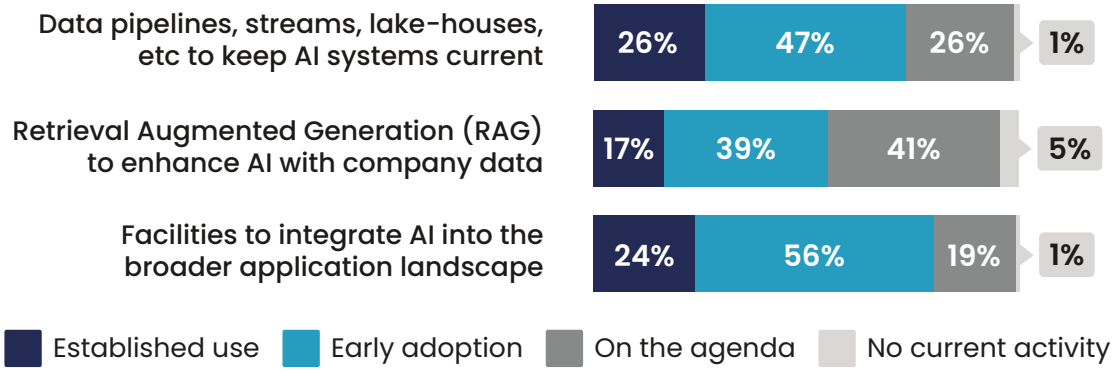

Firstly we need mechanisms in place to source data via pipelines or real-time streams, and where necessary pull it together, e.g. through centralized data lakes or lake houses (which may be physical or virtual/federated).

Source data then needs to be converted into an AI-readable form, via a process known as vectorization. This is where Retrieval Augmented Generation (RAG) solutions and vector stores come into play.

For many use cases, it's then necessary to surface AI solutions as features or capabilities embedded into broader applications, e.g. contact center environments, customer or employee self-service portals, business administration systems, etc. This is where development tools, APIs, etc are important.

Beyond these core capabilities, data-driven AI platforms typically provide a whole range of other facilities to streamline different aspects of implementation and operation.

These are too numerous to cover in detail, but examples include tools to monitor and manage AI models and optimize costs.

Examples of some more advanced or specialist capabilities

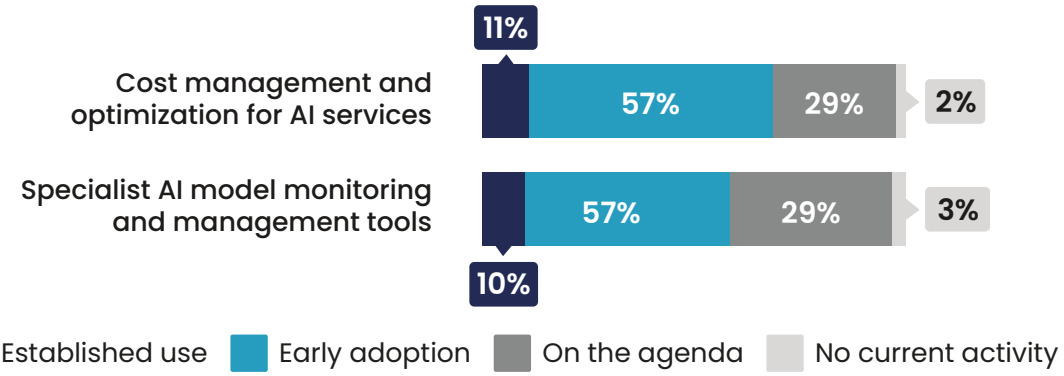

Data readiness must be dealt with proactively, not as an afterthought

Standing back from all of the detail, the most important enabler of success with AI from a strategic perspective is laying the right data and information delivery foundations. In truth, this imperative exists independently of your AI ambitions, plans and activities.

As we explored in a previous research report from 2020 entitled "The Road to Becoming a Data-Driven Business", a strong cultural, procedural and technology foundation for managing and unlocking value from your data delivers a wide range of benefits. These range from increased operational efficiency as a result of data-driven automation, to more effective decisions in the boardroom enabled by real-time business insights.

As part of that earlier study we developed a 'data maturity model' to assess how well organizations were doing, but that was before the upsurge in interest around AI. For this latest study, we reused the essence of that model, but added a series of 'tests' to assess the level of foundational AI capability in place. To put it simply, we adapted the model to take account of the kind of AI platform requirements we have been discussing (see Appendix A for more details).

The result is the model shown below which is based on the principle that a good level of generic data maturity is required, together with reasonable foundational AI capability, for strategic, scalable AI adoption.

As you can see, while it's widely reported that most organizations are active with AI, fewer than a third are laying the foundations for genuine enterprise-class AI initiatives.

Data + AI maturity model

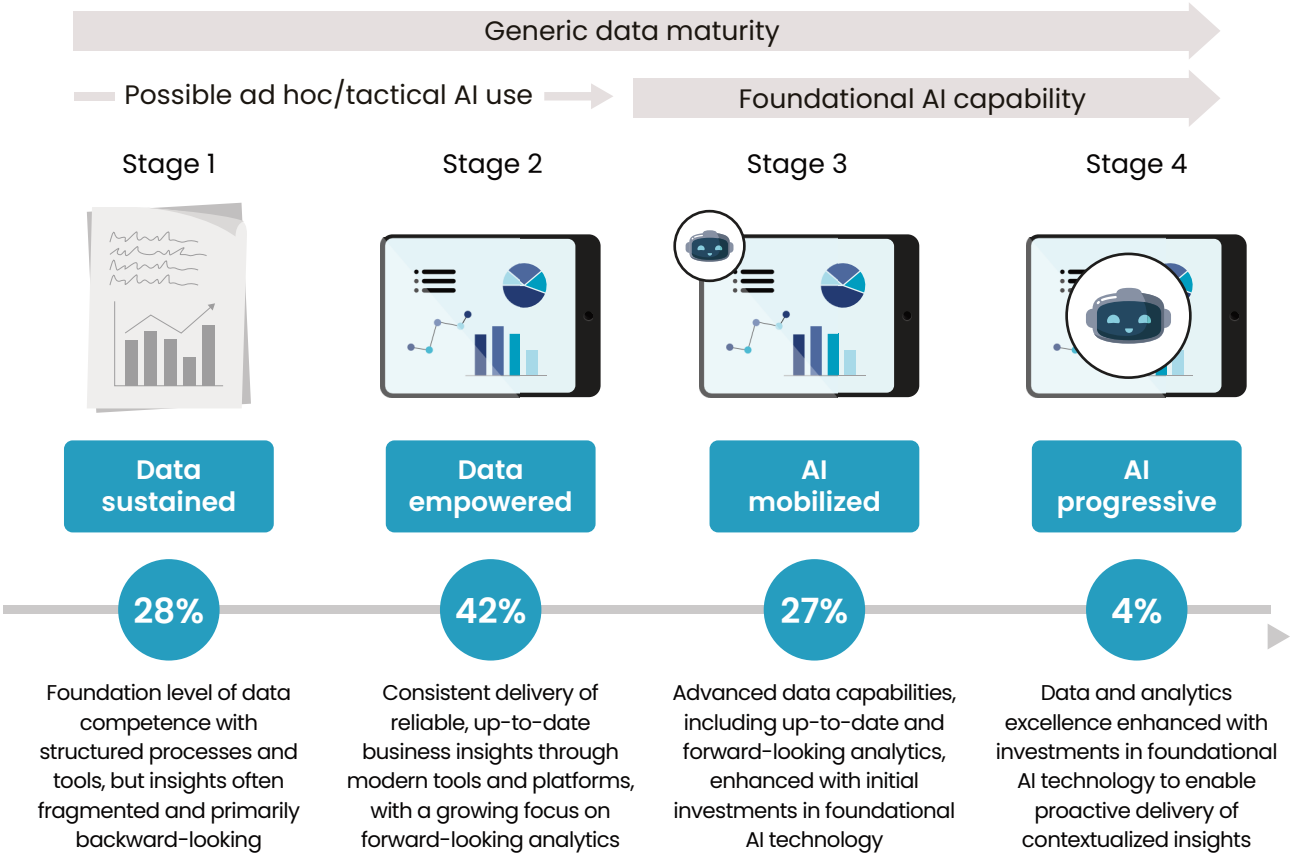

Early movers are setting themselves up for success

An obvious question we need to address is why maturity based on a combined set of data and AI assessments matters - i.e. what are the benefits of achieving a higher maturity level?

In our 2020 study, we saw a strong correlation between data maturity and business level KPIs in areas such as decision-making, productivity, operational efficiency, employee experience and customer satisfaction. In this latest study, however, there were hints of such correlations, but the picture was inconsistent overall.

At one level this was disappointing, but at another it is perfectly understandable as we are very early in this market and initial AI initiatives have not yet had a chance to impact overall organizational performance. That said, some positive indicators did come through quite convincingly.

Benefits of higher Data + AI maturity

- ✓ Implementation of good practices
- ✓ Comfort with different AI deployment types
- ✗ Business performance KPIs (too early to tell)

From our observations (see Appendix B), more mature organizations have a higher level of comfort with different types of AI deployments, consistent with the notion that once you have good foundational capability in place, you can move forward with new and different initiatives more confidently. We've seen this kind of pattern before in other areas, e.g. early movers in the area of hybrid/multi-cloud platforms were able to exploit cloud benefits more aggressively.

Another way in which more mature organizations are seeing benefits is in their ability to implement good practices, e.g. around governance, ethics and training. The principle here is that a robust and consistent foundational approach makes it easier to develop and institutionalize standard practices as you're not having to reinvent policy on a project-by-project basis. It's also worth highlighting that good practices are not just about managing cost and risk. Once they are in place, the organization can move much more quickly on new AI initiatives, accelerating the time to value.

With these factors in mind we anticipate it only being a matter of time before early movers in the AI platform space see their investments and efforts bear fruit in terms of tangible business outcomes.

The Data + AI maturity advantage

Final thoughts: AI as a catalyst for broader system modernization

While the research reported in this document has its roots in data maturity, we deliberately chose to anchor our discussion in AI, as this topic is receiving so much attention right now. This felt very natural, however, given that without good quality enterprise data to provide business specific context and substance, AI represents little more than an entertaining side show.

But if we change perspective and look at the research through a data lens, an interesting opportunity emerges. Beyond creating an investment imperative relating to AI specific platforms, AI initiatives frequently also require the strengthening of generic data management and delivery systems. This might translate to moving from data warehouses to more flexible data lakes or lakehouses, for example, or implementation of real-time data streaming platforms.

Such investments are likely to drive benefits well beyond your AI related activities. In effect, you will be tackling data-related challenges that continue to hold up progress in a whole range of other areas, from routine reporting and analytics, through data-driven process automation, to customer management and much more. It's not that AI won't find it's way into these domains sooner or later - it probably will - but in the meantime, so many systems and processes can benefit right away from better and more timely data.

Such spin-off benefits are worth bearing in mind as you construct business cases for AI initiatives. If you expand your view, AI won't just deliver value directly, it could also act as a catalyst to unlock a broader set of data platform investments that will bring your business even closer to that 'data-driven' Nirvana.

Appendix A
The maturity model in more detail

The Data+AI maturity assessment model was designed to evaluate organizations based on 16 key indicators across four main areas critical to success with enterprise AI initiatives. While each indicator test involved specific assessment criteria, they can be broadly understood from the overview captured in the diagram below.

Data+AI maturity assessment

Culture & mindset

- Attitude to IT investment
- The role of data
- Sharing behavior
- User training/coaching

User experience

- Currency of information
- Access mechanisms
- Consistency of sources
- Reliance on IT team

Data management

- Strategy & governance
- Security/risk responsibilities
- Security/risk measures
- Data quality and integrity

Foundational AI capabilities

- Integrated AI platform(s)
- Data pipelines, streams, etc
- Implementation of RAG
- AI/application Integration

The Culture & mindset tests examine attitudes and behaviors around technology investment, data assets and knowledge sharing. The User experience indicators focus on how well data and insights are delivered to support decision-making. Data management tests assess the robustness of governance, security and quality controls. The Foundational AI capabilities category evaluates readiness in terms of platforms and integration.

Organizations were scored against each indicator and placed into one of four maturity levels based on their aggregate performance: Data sustained, Data empowered, AI mobilized, or AI progressive. These levels reflect increasing sophistication in both general data competence and AI-specific capabilities.

This approach builds on established data maturity assessment methodology while incorporating new tests to account for emerging AI platform requirements. The goal was to provide an objective framework for evaluating organizational readiness for strategic AI adoption.

Appendix B
Notable correlations in the data

The charts below illustrate some notable patterns that emerged when analyzing responses based on Data+AI maturity level.

They illustrate how higher maturity correlates with greater comfort around AI adoption and stronger implementation of good practices.

At a high level, how comfortable are you with the idea of rolling out the following types of AI solutions in your organization?

	AI assistants or copilots				Simple task oriented AI agents			
AI Progressive	53%	41%	6%	0%	59%	41%	0%	0%
AI Mobilized	23%	74%	2%	2%	51%	36%	9%	5%
Data Empowered	21%	70%	7%	2%	47%	40%	10%	4%
Data Sustained	9%	40%	47%	3%	17%	44%	28%	11%

	AI-enhanced business applications				AI agents running on-premises			
AI Progressive	59%	41%	0%	0%	12%	77%	0%	12%
AI Mobilized	26%	67%	5%	3%	10%	68%	9%	13%
Data Empowered	26%	56%	15%	3%	5%	68%	22%	5%
Data Sustained	9%	48%	34%	9%	4%	49%	36%	12%

Extremely comfortable | Somewhat comfortable | Reserving judgment | Not Comfortable

To what degree are copilots and chatbots being used in your organization at the moment?

AI Progressive	35%	41%	18%	6%
AI Mobilized	12%	56%	28%	4%
Data Empowered	8%	67%	20%	5%
Data Sustained	5%	63%	24%	9%

Rollout within a full policy/ governance framework

Proactive rollout with guidance and support

Bottom-up, user-led adoption

No official use (possibly some unofficial)

And what about the use of AI to underpin or enhance custom applications that you develop in house or commission from systems integrators?

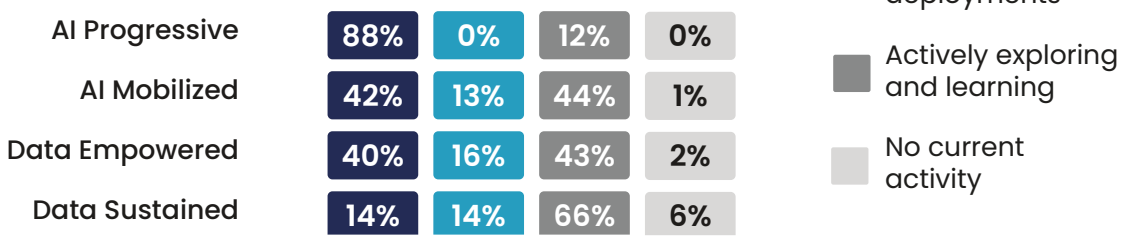

Regarding your AI activities, to what degree have you adopted the following:

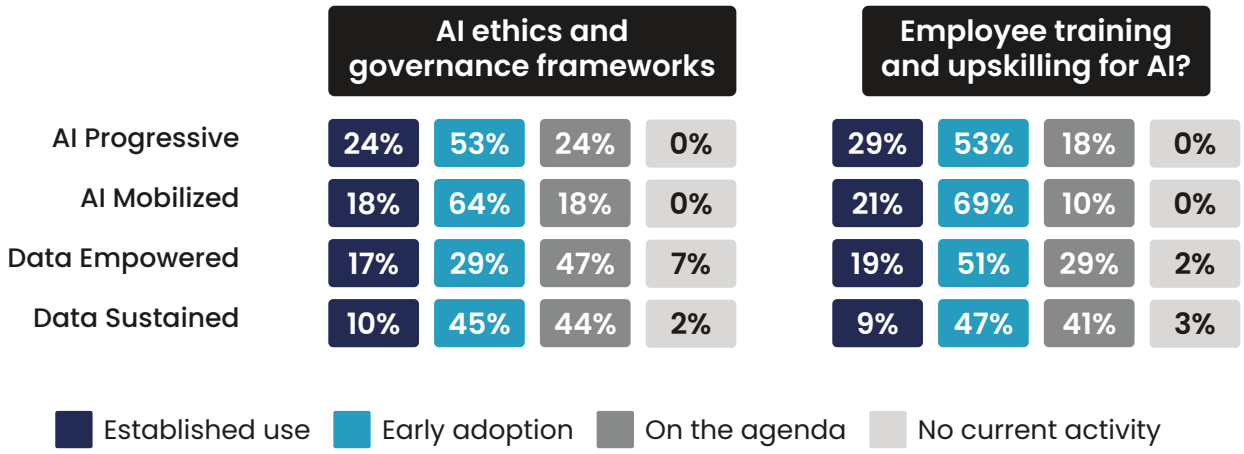

About

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we help busy IT and business professionals get up to speed on the latest technology developments and make better-informed investment decisions.

For more information and access to our library of free research, please visit www.freeformdynamics.com.

About Fujitsu

At Fujitsu, we're passionate about using technology to create a more inclusive, sustainable and trusted future. It drives everything we do. Throughout our history, we've supported businesses and society through delivering robust and reliable IT systems.

Find out more about our business, our history, our philosophy and the countries we operate in, please visit www.fujitsu.com.

Research note

When viewing the charts in this report, please note that percentages may not always sum to exactly 100% due to rounding to the nearest whole number. This is a standard presentation convention and does not affect the validity of the data presented.

Terms of Service