



Couchbase FY 2026 CIO AI Survey:

The Race to Ride the AI Wave

A COUCHBASE RESEARCH REPORT: INVESTIGATING HOW INVESTMENT, EXPERIMENTATION AND DATA STRATEGIES ARE DRIVING AI SUCCESS.

Executive Summary

We live in an AI world. The rise of generative AI (GenAI) is already evolving toward the age of agentic AI, as enterprises seek to use these breakthrough technologies to unlock new opportunities, uncover operational efficiencies and improve their overall business performance.

In this environment, a modern data strategy is more important than ever — an organization's data can make or break its AI ambitions. Couchbase's eighth annual survey of IT leaders investigates how enterprises are riding the AI wave, and whether they can fully take advantage of its capabilities.

Couchbase asked 800 global IT decision-makers from businesses with 1,000+ employees, in sectors including finance, healthcare, gaming and more, about their progress: Is investment still growing? What is the cost of falling behind? Do developers have the freedom to experiment or are rigid restrictions holding them back? Do enterprises understand data enough to truly control AI? Are CIOs still confident in their ambitions?

The study found that **AI investment is growing 52% year-on-year**, faster than digital modernization as a whole. Yet progress is bumpy, with data challenges, and a lack of budget and skills holding businesses back. **Falling behind in AI could result in an annual loss of up to \$87 million for enterprises.**



Solving these challenges will be critical to realizing AI's value, fostering a culture of experimentation and early adoption. Enterprises who encourage experimentation see more consistent AI success: projects are **10% more likely to enter production**, and they experience **13% less wasted AI spend** than enterprises with a more restrictive approach.

Enterprises with a stable data strategy and those that understand the data behind AI have greater control, and are **33% more likely to be prepared for agentic AI**. Although, **21% say their control over AI is insufficient or nonexistent**. Having the right strategy; consolidating AI architecture to reduce complexity, costs and risk; and implementing controls that protect the organization without overly restricting innovation are all key drivers of AI success.

Despite the challenges, IT leaders are highly enthusiastic about AI, with **73% of CIOs excited for its potential**. Organizations must take control of their data and build strong foundations today to turn AI into a competitive advantage.


PART ONE: THE RACE TO RIDE THE AI WAVE

From the earliest days of machine learning and analytics, to the rapid adoption of GenAI and now agentic AI, AI investments are showing no sign of slowing. Almost half of all digital innovation and modernization investment in 2024-25 was in AI (**Fig. 1**), and that figure is accelerating. While overall investment in digital modernization is expected to increase by 35% in 2025-26, AI spend will surge by 51%. (**Fig. 2**)

Fig. 1 – Pull-out statistics


- 45% – proportion of digital spend on AI, 2024-25
- \$30M – average AI spend in large enterprises (10,000+ employees)

Fig. 2 – Expected increase in digital modernization and AI spend, 2025-26



AI budgets are evenly split, and expected to remain so. GenAI, agentic AI and other forms of AI such as machine learning each attract around one-third of total AI spend. **(Fig. 3)** However, this even split masks a crucial truth: the fact that the much newer technologies of GenAI and agentic AI have already reached parity, which suggests extremely rapid adoption.

Fig. 3 – Current and expected investment in different approaches to AI



This rapid evolution opens up new opportunities for organizations, from state-of-the-art applications and services to greatly increased employee effectiveness. Yet it is also a cause of disruption to their environment. Enterprises that cannot keep pace with this evolution could still find themselves falling behind. **(Fig. 4)**

Fig. 4 – Pull-out statistics

- 73% of enterprises say AI is already causing a major transformation and upheaval of the technology environment
- 66% of enterprises are concerned that AI and different approaches to AI are evolving faster than the organization can keep pace.

This transformation and upheaval is a potential catalyst for major change and represents a defining moment for enterprise competitiveness. Organizations struggling to match AI's rapid evolution risk watching nimble startups and forward-thinking competitors capture their market share. Conversely, enterprises that harness this technological disruption can leapfrog established rivals. Despite the pressures, enterprises remain largely optimistic. **(Fig. 5)**

Fig. 5 – Pull-out statistics


- 59% of enterprises are concerned that their business will be displaced by smaller competitors who are better able to take advantage of AI
- But 79% believe that their business can use AI to displace larger, less agile competitors.



PART TWO: THE COST OF FALLING BEHIND AI INNOVATION


Respondents predict that businesses that cannot make use of AI in a timely manner will lose, on average, 8.6% of their revenue. For our sample, that equates to an annual loss of almost \$87 million. **(Fig. 6)**

Fig. 6 – Average loss from inability to make use of AI: \$87 million



Respondents know the clock is ticking. Almost all (96%) say that there is a deadline by which their organization needs to have embraced AI. More than a quarter say it has already passed and 87% say it is within the next six months. **(Fig. 7)** Organizations that miss this deadline could begin seeing the costs above rise even more.

Fig. 7 – Deadlines to embrace AI



The consequences of missing these deadlines extend beyond revenue loss — they threaten fundamental business viability. IT leaders recognize that delayed AI adoption creates a cascading series of competitive disadvantages, from the inability to take advantage of new market opportunities to only surviving in narrow, niche sectors where AI relevance is minimal. **(Fig. 8)**




Fig. 8 – Risks facing businesses that do not take advantage of AI

1. A lack of agility – they will not be able to transform or act as quickly (e.g., taking advantage of new business opportunities)
2. Losing customers as they cannot meet demand and expectations for AI-driven services
3. Losing position to, or being bought out by, competitors who are better able to use the technology
4. Losing employees who expect AI-driven services to improve their working lives
5. Struggling to attract finance or secure an IPO
6. Only surviving in very specific sectors where AI is not relevant

These risks underscore the need for businesses to keep pace with AI developments. The data reveals a market in transition, with encouraging momentum tempered by concerning gaps. While most enterprises are on target with their AI adoption, and almost one-third (30%) are ahead of target, there is a significant minority (18%) who are behind their timelines. **(Fig. 9)**

Fig. 9 – AI adoption progress



AI adoption has proven to be challenging. Almost every organization (99%) has encountered issues, such as with organizational buy-in, accessing/managing data and a lack of skills, that prevented them from pursuing a new AI project or caused an active AI project to suffer setbacks. **(Fig. 10)** Successful AI deployment hinges on organizational buy-in with clear risk management, solid data infrastructure and skilled technical teams.



Fig. 10 – Issues preventing or disrupting AI projects

Issue preventing projects	% affected	Issue disrupting AI projects	% affected
Perception that the risk of failure was or had become too high	45%	Perception that the risk of failure was or had become too high	45%
Problems accessing or managing the required data	42%	Inability to secure the necessary budget or stay within budget	39%
Inability to secure the necessary budget or stay within budget	40%	Lack of confidence that the project would meet security or compliance demands	36%
Lack of buy-in or support from across the organization	33%	A lack of direction from senior leadership on the precise goals of the project	31%
Lack of confidence that the project would meet security or compliance demands	28%	Problems accessing or managing the required data	28%
Lack of skills to deliver the project	25%	Lack of buy-in or support from across the organization	25%
A lack of direction from senior leadership on the precise goals of the project	24%	Lack of skills to deliver the project	21%
Lack of buy-in or support from the C-suite	20%	Lack of buy-in or support from the C-suite	15%

These issues all incur costs. First, there is the proportion of investment wasted on projects that never see the light of day or cannot meet expectations. **(Fig. 11)** More importantly, each failed project delays enterprises' ability to make full use of AI and meet strategic goals. The result is an overall impact on revenue. On average, these issues caused enterprises to delay their strategic goals by almost six months, with a potential cost of up to \$42 million. **(Fig. 12)**

Fig. 11 – Pull-out statistic

- 14% of IT modernization investment spent on prevented or disrupted projects
- 17% of AI investment spent on prevented or disrupted project



Fig. 12 – Average delay in strategic goals and potential cost of delay

Region	Delay (months)	Potential cost (\$ millions)
Total	5.84	42
US	5.50	38
Europe	5.84	42
UK	5.68	28
France	6.24	49
Germany	5.58	52
Turkey	5.18	48
Japan	5.54	45
India	6.03	42
Australia	6.42	43
Singapore	6.38	43



PART THREE: FOSTERING A CULTURE OF EXPERIMENTATION

Most IT decision-makers remain concerned about their organization's AI approach — whether they're losing ground to early adopters, rushing implementation too quickly or failing to capitalize fast enough due to “decision paralysis.” (Fig. 13)

Fig. 13 – Pull-out statistics

- 78% of CIOs say early AI adopters will see by far the greatest success and become industry leaders
- 64% of CIOs are concerned that “decision paralysis” means their organization is not taking advantage of AI as quickly as it could
- 59% of CIOs are concerned their organization is rushing to take advantage of AI without fully understanding its end goals

This is not contradictory thinking: it reflects the rational response of leaders who understand they must adopt AI quickly to stay competitive, yet lack the roadmap to achieve it effectively.

One vital factor in unlocking AI's potential is experimentation. While many CIOs are concerned that projects may increase understanding but not provide concrete results, the vast majority recognize the value of testing new ideas and approaches. (Fig. 14) Experimentation can provide insights and learnings that will make the next project more likely to succeed.

Fig. 14 – Pull-out statistics

- 81% of CIOs agree education and experimentation are critical elements of AI development
- 74% of CIOs say that even failed AI projects have value in terms of the learnings they provide for next time
- 67% of CIOs are concerned that developers are spending time on projects that may increase understanding but do not create tangible results.

Experimentation is one area where the corporate approach to AI can have an outsized influence. At one end of the scale, enterprises can be highly restrictive — not allowing any experimentation and ensuring every project is embarked upon with a clearly defined goal. On the other end, enterprises can encourage experimentation and even make it an expected part of developers' role.

(Fig. 15) Interestingly, these two extremes seem the most successful in terms of AI projects advancing from proof of concept (POC). While on average 38% of projects leave POC and enter production, both are noticeably higher — with a 10% increase on average.




Fig. 15 – Approaches to experimentation

Approach	Proportion of organizations using this approach	Proportion of AI projects that enter production
Restricted – no experimentation, every project has a definite agreed goal	9%	42%
Controlled – experimentation is within strict parameters and must be agreed beforehand	37%	36%
Walled garden – unlimited experimentation within agreed parameters	25%	37%
Open – developers can experiment without limits	9%	41%
Encouraged – developers are expected to spend a certain proportion of their time experimenting	20%	42%

Another benefit of experimentation is the proportion of AI investment wasted on projects that are prevented or disrupted. Enterprises that encourage experimentation report the smallest amount of wasted spend: 3% lower than the average, and 13% lower than enterprises where experimentation is restricted. (Fig. 16)

Fig. 16 – Pull-out statistics

- 13% lower wasted AI spend in enterprises where experimentation is encouraged

The difference may only be a few percentage points, but spread across entire AI budgets and the potential impact on revenue, the consequences are significant. This suggests that encouraging experimentation, and identifying and potentially solving issues before they can prevent or disrupt projects, can produce notable benefits for enterprises and increase the chance that AI projects will be successful and profitable.



PART FOUR: UNDERSTANDING DATA IS KEY TO KEEPING AI UNDER CONTROL

Without a deep understanding of data, AI strategies cannot succeed. Poor quality data can give rise to hallucinations or introduce biases. If data isn't recorded and accessible in real time, AI will make conclusions based on outdated information or be unable to share timely advice. And, if the organization doesn't have effective data governance in place, AI can easily become a security and compliance risk. Every enterprise has encountered AI issues that stem from a lack of data control. **(Fig. 17)**

Fig. 17 – Data issues encountered by enterprises

AI sharing another organization's IP or other proprietary information as part of an answer	47%
AI using or accessing proprietary data from within the organization	43%
AI "hallucinations" that employees acted on before they recognized them as such	39%
AI "hallucinations" that employees could identify and avoid acting on	34%
AI operating in a way that works contrary to established best practices	29%
Losing the efficiency benefits of AI through having to double-check its conclusions	23%

Beyond this, respondents recognize that data control forms the foundation of successful AI use. Whether improving control to bring applications closer to end users or identifying where data is at risk, organizations understand that current approaches fall short of AI's demands. **(Fig. 18)**

Fig. 18 – Pull-out statistics


- 77% of enterprises need to bring critical business applications closer to the end user at massive scale, and say better control of data is the only way to enable this
- 72% of enterprises' understanding and control of data needs to be exponentially higher than before to use AI effectively and safely
- 62% of enterprises do not fully understand where they are at risk from AI – e.g., through security or data management issues

The stakes couldn't be higher. AI represents a generational shift with considerably more potential outputs and variables than previous technologies, demanding greater levels of enterprise data control and governance.

Yet organizations aren't necessarily meeting the challenge. While respondents are somewhat confident they understand AI, they are less certain that they understand the data required to power AI and other applications. A striking 70% of respondents say their understanding of the data needed to power AI is, at best, "incomplete," revealing sizable knowledge gaps. At the opposite extreme, only 2% claim "complete" understanding, which is a sobering reality considering AI's data-centric nature. **(Fig. 19)**



Fig. 19 – Enterprises' understanding of AI and the data behind it



As we can see, there is a concerning trend: enterprises' claimed understanding of AI is greater than their understanding of the data behind it. This means a minority of organizations believe they fully understand the AI systems they are using.


This lack of understanding will make it harder to control AI and use it effectively. It also makes it difficult to know when using AI is appropriate and to ensure everyone in the organization has an understanding of different AI approaches. The result can be a lack of control over AI use and development, including not understanding what progress the business is making. (Fig. 20)

Fig. 20 – Pull-out statistics

- 71% of organizations believe they have used GenAI to solve an issue when another approach might have been simpler or more effective
- 66% of CIOs are concerned that not everyone in their organization who needs to, understands the differences between different AI technologies, their strengths and weaknesses

The opposite is also true. Organizations with a greater understanding of data are less likely to have stakeholders that do not understand the nuances of AI, and can better recognize how new technology puts them at risk. Additionally, they are more likely and better prepared to use different forms of AI. (Fig. 21)

Fig. 21 – Preparedness for using AI



To gain a better understanding of AI preparedness, enterprises need to focus on data management. There are multiple factors that form a successful data management strategy, including control over data protection, storage, access and use, and the technology to manage and access data at scale and at speed. (Fig. 22)

Fig. 22 – Data management practices in place at enterprises

- 49% of enterprises have complete control over where data is stored, who has access and how it is used
- 44% have the ability to access, share and use data with minimal latency
- 39% have the tools that prevent proprietary data being shared outside the organization
- 36% of organizations have a high-performance database that can manage unstructured data at high speed
- 35% have clear and thorough best practices that allow developers to use data safely and effectively
- 28% have the ability to scale data processing to meet immediate needs without unnecessary spending
- 25% have a consolidated database architecture so that AI applications cannot access multiple versions of data
- 21% have the ability to perform real-time analytics on large amounts of data
- 16% have a vector database that can store, manage and index high-dimensional vector data efficiently

PART FIVE: UNDERSTANDING REQUIREMENTS AROUND AI DATA GOVERNANCE

The imperative to manage and understand data stems from the fundamental recognition that AI requires robust control via governance and frameworks.

(Fig. 23) Without this control, organizations face a potentially insurmountable dilemma. There are significant risks in sharing proprietary corporate data with AI, which could eliminate the competitive advantage that data provides. Yet this data is essential for combating AI hallucinations and improving its contextual accuracy. This governance and control are vital to providing the guardrails necessary to address this dilemma.

Fig. 23 – Pull-out statistics

In order to succeed with AI:

- 79% of CIOs need strict governance in place
- 78% need strict control over data in place



Effective control must address the multifaceted challenges that threaten safe, strategic AI implementation. Organizations should ensure AI is being used in the proper way, by the appropriate users and with the right data and clear boundaries in place. Businesses should protect against data exposure and maintain user competency in AI capabilities and limitations. (Fig. 24)

Fig. 24 – Top challenges to effective, safe use of AI in-house

1. Monitoring and managing GenAI application use
2. Sharing and accessing data quickly enough to ensure peak performance
3. Encouraging safe experimentation that will result in increased understanding and new capabilities
4. Preventing inadvertent IP theft or security issues
5. Setting realistic goals and expectations of what the technology can do
6. Ensuring access to computing and storage resources
7. Ensuring architecture is high-performance and flexible enough to support GenAI
8. Effective data management
9. Training end users

These challenges reinforce themes throughout this study, but they converge on one important insight: architectural simplicity enables better governance. When examining requirements for safe experimentation, data control and user access to essential tools, complexity emerges as the primary obstacle to effective AI management.

The more complex AI architecture is, the less governable it becomes, making it harder to control. At a minimum, ensuring visibility and control over sprawling technology stacks requires more resources. In the most severe cases, uncontrolled architectural sprawl exposes enterprises to unacceptable risk.

Because of this, consolidation is a priority for enterprises. Every respondent surveyed has reported that their organization has taken action to reduce the size of their AI stack, with clear consensus supporting this direction. (Fig. 25)

Fig. 25 – Pull-out statistics

Why consolidate the AI stack?

- 75% of CIOs believe the current move to AI is an excellent opportunity to consolidate and simplify technology stacks in general
- 74% would rather have fewer, higher-performance data platforms that can manage database, analytics, AI and other services than a higher number of more specialist platforms
- 69% believe AI will be easier to control if there are fewer technologies in the AI stack
- 63% believe reducing the size and scale of the AI stack is essential to addressing issues around energy use and emissions



Consolidation is underway, with more than a third (36%) of enterprises reducing the number of data platforms used internally. These organizations recognize that modern, AI-ready data platforms lessen the need for legacy or highly specialized platforms that, despite performing essential functions, lack the adaptability to support new technologies. (Fig. 26)

Fig. 26 – Actions taken to consolidate the AI stack

Updating infrastructure strategy to consolidate the amount of IT infrastructure used	51%
Using AI applications to perform multiple functions and replace legacy applications	45%
Reducing and consolidating the number of third-party AI services and applications in use	42%
Reducing the number of data platforms used in the organization	36%
Limiting AI use to certain areas of the business	30%
Reducing the number of AI projects the organization begins	23%

Simultaneously, enterprises will continue to review and modernize their IT architecture as demands on databases, data centers and people grow. Inaction is not an option: on average, enterprises say their current architecture has an average lifespan of 18 months before it can no longer support in-house AI applications. (Fig. 27)


Fig. 27 – Pull-out statistic

- 18 months – average lifespan of organizations' current AI architecture

There are multiple approaches enterprises are already taking to improve control and reduce the risks involved in AI use. (Fig. 28) Each of these is valuable, but most should also form part of a comprehensive data management strategy.

Fig. 28 – Actions taken to control AI and AI use

Controlling AI access to data	50%
Controlling employee access to AI applications	43%
Investment in company-wide training on correct AI use	38%
Putting guardrails in place to ensure AI complies with regulations	32%
Creating revised AI usage policies	31%
Putting in place disciplinary procedures for improper AI use	25%
Banning the use of third-party AI applications that share and rely on public data	22%
Deleting duplicate data to reduce the risk of hallucinations	17%



An important aspect is striking the right balance between being overly restrictive and overly permissive. As we have seen with experimentation, an overly restrictive approach can stifle innovation and slow progress. And a highly permissive approach could expose the business to unnecessary risk.

Fortunately, this is one area where most respondents are broadly satisfied with their organization's level of control. While there is room for improvement, only 7% believe that the level of control within their organization is overly restrictive. A concerning 21% say their enterprise has insufficient, or even zero, control over AI use. **(Fig. 29)**

Fig. 29 – Pull-out statistics

How do CIOs feel about their organization's level of control over AI?

- 45% Satisfactory – it reduces some risk and allows employees freedom, but could be improved
- 28% Perfect
- 20% Insufficient – it allows employees too much freedom and increases risks
- 7% Restrictive or too stringent – it reduces risks but doesn't maximize benefits
- 1% Nonexistent – zero control over AI use

For this vulnerable 21%, establishing robust controls represent the difference between realizing AI's full potential and falling behind as new AI capabilities are introduced. This is especially important as enterprises anticipate breakthrough opportunities from emerging technologies like agentic AI. **(Fig. 30)**

Fig. 30 – What most excites enterprises about the potential of agentic AI?

1. Rapid prototyping and testing of new ideas – e.g., for business strategy, marketing and sales materials, product designs
2. Improving customer experiences to meet or exceed expectations
3. Identifying new business trends and capitalizing on them quickly
4. Identifying and combating rapidly evolving security threats
5. Maintaining parity with competitors who have invested or are investing in agentic AI
6. Making employees in other business units more efficient by performing manual tasks for them
7. Increasing developer productivity

Ultimately, CIOs know that AI defines the future, and that the right level of control — over users, data and architecture — is crucial to its success. Most respondents are clear that not using AI is the greatest risk of all. And there is a clear compulsion to keep using and experimenting with the technology to uncover new, exciting value and succeed in an AI-driven world. **(Fig. 31)**



Fig. 31 – Pull-out statistics

CIOs' last words on AI:

- 73% say “there are challenges involved, but I am still excited about AI’s potential”
- 68% say “I feel compelled to use AI more”
- 64% say “not using AI at all is a far greater risk than any of those involved in implementing it”

CONCLUSION

While AI adoption challenges can overwhelm organizations, CIOs remain clear-eyed about both the opportunity and path forward.

The reality is stark: with 80% of CIOs agreeing that we are entering an AI world, enterprises have no choice but to build applications for it. Creating and operating innovative AI applications, at scale, will be one of the hallmarks of a successful enterprise.

In this environment, success hinges on the right approach to data management. Enterprises must implement robust controls and simplify architectures using comprehensive and unified multipurpose data platforms that support the variety of data involved (structured, unstructured, vectors, prompt inputs, textual responses from genAI, etc...) to provide the exact, consistent capabilities they need throughout every GenAI interaction.

With proper foundations in place, organizations can confidently navigate their AI environment and the data that supports it, unleash AI’s full potential without compromising business security, encourage productive experimentation and position themselves to capture maximum value from their AI investments and transformation.

METHODOLOGY

The report is based on an online survey conducted in April 2025 by [Coleman Parkes](#), an independent market research organization, of 800 senior IT decision-makers, such as CIOs, CDOs and CTOs, in organizations with 1,000 employees or more in the U.S., U.K., France, Germany, Turkey, Japan, India, Australia and Singapore. IT leaders in industries across financial services, retail, manufacturing, telecommunications, healthcare, energy and utilities, gaming, and travel and hospitality were surveyed.

Modern customer experiences need a flexible database platform that can power applications spanning from cloud to edge and everything in between. Couchbase’s mission is to simplify how developers and architects develop, deploy and run modern applications wherever they are. We have reimagined the database with our fast, flexible and affordable cloud database platform Capella, allowing organizations to quickly build applications that deliver premium experiences to their customers – all with best-in-class price performance. More than 30% of the Fortune 100 trust Couchbase to power their modern applications. For more information, [visit **www.couchbase.com**](https://www.couchbase.com) and follow us on X (formerly Twitter) @couchbase.

