



CLOUDERA


SURVEY REPORT

The Great AI Re-Architecture

Why the next era of AI depends
on data architecture

Table of Contents

Executive Summary	3
The Great AI Re-Architecture	4
An AI-Driven Data Architecture Reset	5
The Path to Scalable AI Runs Through Effective Governance	6
The Emerging Hybrid AI Landscape	7
The Great AI Re-Architecture: A Global View	8
What Comes Next? Delivering AI Anywhere	8
Methodology	9
About Cloudera	10



Executive Summary

It's no secret that AI is reshaping how enterprises approach nearly every facet of operations. What started in small, siloed, pilot projects has quickly exploded into a full redefining of the enterprise landscape. Businesses aren't just looking for shiny new technologies to show off; they are looking for solutions that deliver real business value and tangible results, quickly and securely.


And it's that inflection point that's bringing enterprise leaders to the same conclusion—the data architectures that have long been the standard in enterprise IT are no longer enough for the governance, performance, cost, sovereignty, and security demands of modern AI.

It's a massive shift: a Great AI Re-Architecture. One where enterprises are shifting to data architectures that can bring trusted AI to trusted data, wherever it lives. And as this re-architecture continues, data movement between on-premises, private cloud, public cloud, edge, and hybrid environments is redefining data management.


So, the Great AI Re-Architecture has arrived and it is not optional, but what does it actually look like? To better understand how organizations are navigating an AI-driven future and its impact on data infrastructures, Cloudera, in partnership with Wakefield Research, commissioned "The Great AI Re-Architecture," surveying 1,500 Enterprise Architects, Cloud Infrastructure Leads, and Data Architects.

The Great AI Re-Architecture

This survey reveals a clear picture of the current AI infrastructure landscape. According to Cloudera's report, 77% of organizations are actively using AI in some form. But at the same time, the path forward to achieving scalability and long-term success is still complicated. In fact, **95% of respondents said they have delayed or canceled projects in the last 12 months entirely due to data governance, compliance, or regulatory issues.** And more than half (55%) reported canceling or delaying more than six projects in that time.



But governance challenges are only one piece of the re-architecture puzzle. Respondents understand that changes are happening and that they are far overdue. The survey found that **75% of respondents said AI integrations had changed their organization's data storage and architecture practices.** At the same time, 72% said their current data architecture requires a significant overhaul to meet their AI goals.



Ultimately, when we talk about a Great AI Re-Architecture, the meaning goes much deeper than modernization. It is about rethinking where data lives, where AI workloads run, how governance travels across environments, and how organizations maintain control over sensitive data while scaling AI. All that, while considering cost controls and flexibility. For many enterprises, that means moving toward hybrid architectures that combine public cloud, private cloud, on-premises, sovereign and edge environments.

The core challenge for enterprise IT leaders will be building a data foundation that delivers trusted AI outcomes at scale.

In this report, we will explore the current state of data architectures and how they are grappling with AI workloads, the evolving governance landscape, and where hybrid approaches enter the architecture picture.

What is the Great AI Re-Architecture

A strategic redesign of enterprise data architecture to more effectively scale AI. This shift is driving leaders to rethink where data resides, where workloads run, and how sensitive data is managed across environments.

An AI-Driven Data Architecture Reset

For many businesses, this is a shift that's already well underway. More than one quarter (26%) of respondents said they were highly AI-driven, with strong leadership support and a culture of experimentation.

Add to that the fact that 77% said they are actively using AI in some form, and it's clear the days of the simple pilot project are long gone. Now it's an issue of scale. With AI becoming embedded across more teams and more core operations, an inadequate data architecture has proven to be a major limiting factor.


What otherwise worked for those isolated pilots is often not enough for producing AI outcomes at enterprise scale—a setting where models require consistent access to governed, high-quality data. This is one of the biggest impacts that the rise of AI has brought to the data architecture space.

The vast majority (84%) of respondents said that AI workloads have caused infrastructure costs to increase. And organizations are already feeling the pressure to overhaul how they store, structure, and manage data. Three-quarters of IT leaders surveyed said AI integrations have changed their data storage and architecture practices, including where data is stored and how it is managed across environments.


84%

said that AI workloads have caused infrastructure costs to increase.


But the driving force for that change varies from place to place. Asked to identify the primary driver behind changes to their data storage and architecture practices, 42% of respondents identified data security, governance, and compliance requirements. The next closest factors cited the need to improve performance/reduce latency (35%), the need for real-time or edge-based AI capabilities (35%), and scaling AI initiatives across the business (33%).



identified data security, governance, and compliance requirements.



cited the need to improve performance/reduce latency.



need for real-time or edge-based AI capabilities.

Other drivers of storage and architecture change included reducing reliance on a single cloud provider (33%), modernizing legacy infrastructure (30%), competitive pressure/industry trends (28%), and the need to reduce costs (25%). In other words, traditional storage capabilities are not enough for AI agentic workloads.

Whatever the driving force may be, a majority (72%) of organizations agree that their current data architectures are due for a significant overhaul to meet AI goals. Even if enterprises are ready to use AI, many are coming to the realization that the foundational infrastructure it relies on is not.

The Path to Scalable AI Runs Through Effective Governance

The Great AI Re-Architecture has ramifications for more than just IT infrastructure. One is the way organizations approach data governance—something that we've already seen ranks highly among drivers for overhauling data storage and architecture. As organizations aim to increase AI adoption, the path to scalability runs directly through data governance. In other words, it can no longer be considered a compliance checklist item to mark off.

Failing to get data governance in order has real consequences—something that many respondents have reported experiencing themselves. Ninety-five percent of organizations said planned AI projects were delayed or canceled in the past 12 months due to data governance, compliance, or regulatory complexity. And more than half (55%) said that they had to delay or cancel six or more projects in the last 12 months.

55%


of respondents said that they had to delay or cancel six or more projects in the last 12 months.

The scale of that disruption reflects how much AI has changed the governance equation. AI systems depend on large volumes of data that may be stored across cloud, on-premises, SaaS, and edge environments. As models draw from more systems and workloads spread across more locations, it becomes harder to maintain consistent policies around access, usage, privacy, quality, and compliance. Enterprises had good governance systems when humans were the only ones accessing data manually. But when thousands of agents provide support to employees, customers or partners the story changes and governance requires another dimension.


Survey respondents are keenly aware of that complexity, (29%) strongly agree that AI integration has made data governance more complex and difficult to maintain. This number sheds insight into where the actual challenge with governance lies. It's less about the regulations organizations are trying to comply with and more about AI systems themselves transforming the enterprise data estate, making it much harder to control.

The survey results show that the challenge is not limited to one particular corner of the data estate or one environment. When asked to rank the environments based on which is the hardest to govern for AI workloads, the largest share of respondents identified private cloud (29%), followed closely by hybrid (24%), public cloud (20%), and the edge (14%). Fewer than one-in-ten (9%) said on-premises environments were the hardest to govern.

Respondents identified when asked to rank the environments based on which is the hardest to govern for AI workloads.



But data isn't stationary. Nearly all (97%) respondents said they move data between environments monthly or more often. And almost one-third (31%) said they move data between environments daily.



This is also what makes private AI and data sovereignty so important. As organizations scale AI across regulated or sensitive data environments, they need the ability to keep data under control while applying consistent policies across the organization, all while ensuring their AI systems operate within the security, privacy, and compliance boundaries that the business requires.

Ultimately, the position of governance has changed. Think less of compliance to-do list items and more of critical infrastructure. Organizations cannot scale AI by adding more models, more compute, or more applications. Enabling AI at scale starts with a unified data foundation that ensures control over 100% of organizational data, wherever it resides.

The Emerging Hybrid AI Landscape

Addressing the governance puzzle of scaling AI becomes even more urgent as organizations move toward more distributed AI architectures. One notable trend is the resurgence of on-premises and private cloud environments. Roughly two-thirds (66%) of respondents said they had moved AI workloads from public cloud back to on-premises or private cloud environments in the last 12 months.

66%

of respondents said they have moved AI workloads from public cloud back to on-premises or private cloud environments in the last 12 months.

That's not to say that public cloud is falling out of favor for enterprises grappling with an AI re-architecture. What it does mean is that IT leaders are putting more emphasis on moving workloads to the environments that they are best suited for—considerations that span performance, cost, latency, governance, availability, and accessibility all sit at the forefront of every organization's strategic evolution.


And as organizations grapple with how best to manage AI workloads, hybrid architecture has only become more prominent in the AI re-architecture that is already underway. Asked where AI inference workloads are running in their organizations, hybrid (31%) ranked as the top option. Though this was followed closely by public clouds (30%) and private clouds (22%), with just 8% pointing to on-premises. Ultimately, organizations are not leaning toward any one environment for their AI inferencing needs. As optimization remains the priority, organizational leaders are actively trying to restructure their data estates to meet those needs. In addition to this, many European customers are starting to consider a new breed called "Sovereign Cloud".

Which is an evolution of the local data center with the scalability of the public cloud and a 100% owned / sovereign footprint.

So, what's influencing the shift in inferencing workloads? Asked which environment currently delivers the best performance for AI workloads, respondents pointed to public clouds (30%) the most, with hybrid (29%) in near-lockstep and private cloud (23%) and on-premises (9%) following behind. This paints a clearer picture: the best approach is not tied to one specific environment.

Organizations want, and need, to bring AI to their data, wherever it lives. Now, with the hybrid advantage coming into form, what does this mean for the future plans of IT leaders? Over the next two years, respondents indicated that they plan to increase cloud spend (29%) the most. But compared with other data architectures, the results show that the future is much more mixed—one quarter (25%) said they plan to place more emphasis on a hybrid-first approach, while nearly another quarter (24%) said they are actually planning to increase on-premises spend and 22% plan to increase edge spend as well.

Over the next 2 years, how do you anticipate your organization will change its data architecture to accommodate AI workloads?




What these numbers show is a clear future where the dominant architectures—those that will allow AI workloads to accelerate, scale, and thrive—will be a hybrid mixture. Organizations will prioritize moving data into multiple environments to get the most out of AI models. Ultimately, it's a shift that requires building flexibility to run the right workload in the right place while maintaining visibility, governance, performance, and cost control across the entire data landscape.

The Great AI Re-Architecture: A Global View

The results make clear that the current AI re-architecture is a global phenomenon, spanning countless regions and countries. But the shifts, priorities, and challenges that underpin those changes reveal some variance. For this report, we will examine these differences across three primary regions – the Americas, EMEA (Europe, the Middle East, and Africa), and APAC (Asia-Pacific).

Does your current data architecture require a significant overhaul to meet future requirements for AI?




Across regions, organizations are consistently entering similar stages of AI adoption. However, respondents in the EMEA region signaled the deepest level of adoption, with 47% saying they use AI widely across their organizations and core operations. In comparison, this number was 36% in the Americas and 30% in the APAC region. Regardless, 100% of organizations across regions reported some level of AI use across the board.

Looking at how data is stored across organizations, additional differences begin to emerge. Respondents in the EMEA region reported the highest use of public cloud (38%), while respondents in the Americas cited public cloud the least (28%). In the case of private cloud, EMEA respondents had the lowest score at just 19%, while those in the Americas cited its use the most at 27%. Organizations in the APAC region followed closely behind at 26%.

As data architectures undergo this transformation, the governance concerns of global organizations reveal some separation in priorities as well as actual impact. Comparing across regions, EMEA respondents reported the highest rate of delaying or canceling at least 6 AI projects in the last 12 months due to data governance, compliance, or regulatory issues at 60%, with respondents in the Americas reporting this at 54% and APAC at 50%. Regardless of the differences, this points to a clearly universal issue as every region eclipsed the 50% mark.

How many AI projects has your organization delayed or canceled in the last 12 months due to data governance, compliance, or regulatory issues?



Interestingly, there is a sizable gap in opinion when it comes to governance by region. Asked to rate the accuracy of the statement “AI integration makes data governance more complex and difficult to maintain,” respondents in the Americas had the strongest reaction, with 81% agreeing, while 65% in EMEA agreed, and 61% in APAC agreed as well.

Taken together, the findings aren’t pointing to separate futures but to different paths toward the same destination. With that being a future driven by the need to deliver on the promise of AI, a task that starts by building an architecture that ensures control and access over data, no matter where it lives.

What Comes Next? Delivering AI Anywhere

The next evolution in the AI journey is here. Gone are the days of pilot projects, experimentation, and limited rollouts. Organizations are pushing to integrate AI and Agentic intelligence deeper and deeper into the core of their operations. But as the integration spreads, more projects are failing, governance practices are revealing critical gaps, and workloads are increasingly on the move from one environment to another. From the governance projects to agentic harnessing and controls. Token consumption becomes another variable in the cost together with compute or storage units.

In short, IT leaders are coming up against a realization: traditional data architectures are in serious need of an update.

And so, the Great AI Re-Architecture is upon us. But it's so much more than just an infrastructure refresh or reset. It's a fundamental shift in how organizations think about data as the foundation for AI-driven operations and its orchestration. The reality that widespread AI adoption creates is one where organizations must position themselves to operate across hybrid environments, across all environments at the same time, protect private and sovereign data, and maintain governance wherever workloads run.

The enterprises that move fastest will be those that can bring AI to their data wherever it resides, use trusted data to power trusted outcomes, and build the flexibility to scale across environments without losing control.

Learn more about how [Cloudera](#) is helping organizations navigate evolving data architecture demands and enabling AI, anywhere.


Methodology

The survey, commissioned by Cloudera and fielded by Wakefield Research (www.wakefieldresearch.com) among 1,500 Enterprise Architects, Cloud Infrastructure Leads, and Data Architects at companies with a minimum size requirement of 1,000 employees in all markets except Spain, Singapore, and S. Africa, for which the company size minimum is 250 employees. The research was conducted in 3 Regions and 9 markets: Americas (U.S. (n600), Canada (n100), Brazil (n100)), EMEA (S. Africa (n100), Spain (n100), U.K. (n200)), APAC (Singapore (n100), India (n100), Japan (n100)) between June 5th and June 22nd, 2026, using an email invitation and an online survey.

About Cloudera

Cloudera is the only hybrid data and AI platform company that large organizations trust to bring AI to their data anywhere it lives. Unlike other providers, Cloudera delivers a consistent cloud experience that converges public clouds, on-prem data centers, and the edge, leveraging a proven open-source foundation. As the pioneer in big data, Cloudera empowers businesses to apply AI and assert control over 100% of their data, in all forms, improving security, governance, and real-time and predictive insights. The world's largest brands across all industries rely on Cloudera to transform decision-making and ultimately boost bottom lines, safeguard against threats, and save lives.

To learn more, visit [Cloudera.com](https://cloudera.com) and follow us on [LinkedIn](#) and [X](#).



CLUDERA

Cloudera, Inc. | 6220 America Center Dr, 5th Fl, San Jose, CA 95002 USA | cloudera.com

© 2026 Cloudera, Inc. All rights reserved. Cloudera and associated marks are trademarks or registered trademarks of Cloudera, Inc. All other company and product names may be trademarks of their respective owners. 9512-01 July 16, 2026