NTT DATA

May 2026

2026 Global AI Report: A Playbook for Private and Sovereign AI

How organizations are designing AI for trust, control and resilience

Ignite tomorrow today.

Contents

Executive summary

As organizations scale AI, they are running into new constraints. Some are internal, including the need to run AI in **private environments** to ensure control, security and performance. Others are external, such as the rise of **sovereign requirements** that influence where data and AI systems reside and operate, and under what conditions.

Almost everyone is experiencing this shift. In fact, **95% of organizations** say private and sovereign AI are important — but far fewer are ready to act on it.

What is emerging is a clear split. Some organizations are redesigning how their AI works: where it runs, how it is governed, and how private and sovereign requirements are built into the architecture from the start. Others are still trying to layer AI into systems that were not designed for that level of control, locality or data-flow constraint.

The difference between the two is starting to show.

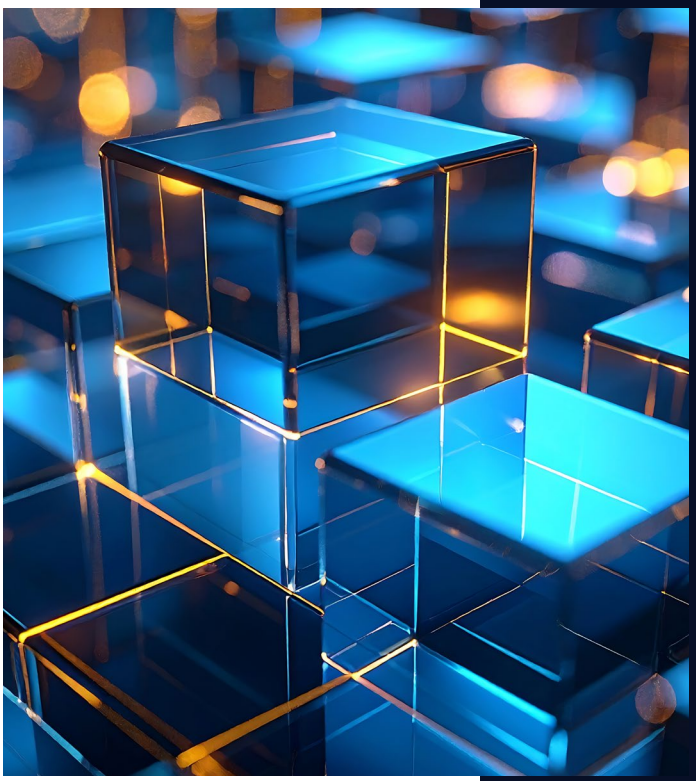

Who are the AI leaders?

Respondent organizations were classified as AI leaders[1] if their survey responses indicated their AI strategy is well-defined or in progress; their level of AI maturity is mature or evolved; and they have realized significantly higher profits from AI than their peers.

These AI leaders are growth and margin outliers. They are:

- Nearly **2.5 times** more likely to post revenue growth of more than **10%**
- **3.6 times** more likely to run at margins of **15%** or more.

They appear across revenue and organization size ranges. Performance advantage is linked to how these organizations design and operate AI.

[1] NTT DATA's 2026 Global AI Report: A Playbook for AI Leaders, the first publication in this series, presents nine characteristics of these AI leaders in detail.

1

AI is running into a wall — and it's not the model

For years, progress in AI meant better models. That is no longer the main issue.

The real friction is underneath, in infrastructure. Running AI in **private environments** already demands greater control over computing, data access and security. Adding **sovereign requirements** introduces another layer, with localization, jurisdictional boundaries and tighter constraints on how data and workloads are handled.

What is becoming clear is that **AI is forcing a redesign of enterprise architecture and infrastructure**. Systems built for centralized, borderless data flows are struggling to support AI that must run in controlled, and increasingly localized, environments. This challenge is widely recognized. About **35% of CAIOs** identify enabling private and sovereign AI as their top barrier to adoption, often requiring significant changes to their existing infrastructure.

The gap between what AI now requires and what existing infrastructure can deliver is widening.

2

Data jurisdiction is becoming an architectural constraint

One of the biggest pressure points is data. In **private AI**, the focus is on controlling access — ensuring sensitive data stays within enterprise boundaries. **Sovereign AI goes further**, placing constraints on where data can physically reside, how it can move across regions and which environments can process it.

These restrictions are not new. Regulations have governed cross-border data flows for years. What is changing is the role data plays in AI. AI systems depend on large, diverse and often geographically distributed data. They rely on continuous access to that data, as well as its movement and recombination, whether for training, fine-tuning or real-time use. In many cases, data can still move legally but not with the speed, scale and fluidity that AI architectures often assume. These considerations are turning what was once primarily a legal issue into a **core architectural constraint**.

Most AI leaders — **nearly 60%** — already cite cross-border data restrictions as a major challenge. This is forcing a rethink of some basic assumptions: where data lives, where models run and how the two come together. As a result, **AI is moving from globally scaled systems to more fragmented, regionally bounded ones**. That introduces trade-offs — between performance and compliance, and between efficiency and control — that most organizations have not had to manage before.

3

Everyone sees the shift, but few are acting on it

There is little debate about the importance of private and sovereign AI. Both are firmly on the agenda. But when we look at what organizations are actually doing, the picture changes. Even though nearly all organizations recognize the importance of private and sovereign AI, **only about one-third (29%)** are prioritizing sovereign AI in a concrete, near-term manner.

This suggests that although organizations understand the need to operate AI in controlled and compliant environments, many are still in the early stages of translating that into scalable architecture and operating models, particularly where sovereign requirements add complexity. What stands out are gaps in both action and confidence. For example, **only 38% of organizations** say they feel highly confident in their cloud security posture — a critical foundation for private and sovereign AI.

The direction is clear, but the follow-through is uneven.

4

Leaders redesign early and move decisively, creating competitive divergence

The organizations that are pulling ahead are not necessarily dealing with fewer constraints; rather, they are responding to them differently. They treat **private and sovereign AI as a foundation** and **a design principle**, aligning infrastructure, governance and operating models early. That is evident from how quickly they can move from pilot projects to scaled deployments, even in complex and regulated environments.

Others are taking a more incremental path, often trying to extend existing setups into sovereign contexts — but this approach is starting to create friction.

What is emerging is not just a technology gap but also a structural one. Many organizations are not set up architecturally or operationally to handle sovereign requirements at scale. This difference is visible in the data, where AI-leading organizations are consistently ahead by **10 to 11 percentage points** in prioritization and readiness indicators.

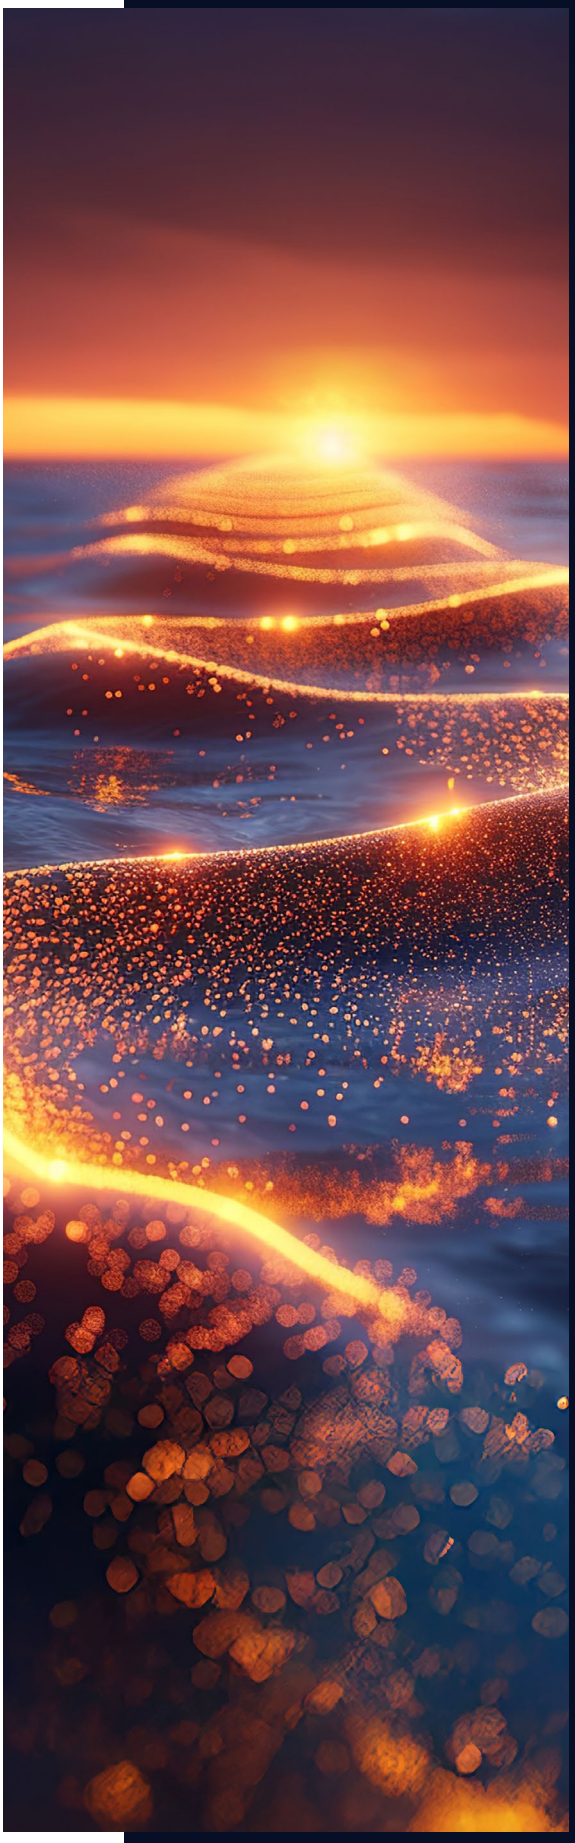

5

Private and sovereign AI sound like independence but are built on tightly orchestrated ecosystems

Both private and sovereign AI increase complexity, but in different ways. **Private AI** concentrates complexity within the enterprise: securing environments, integrating systems and managing internal data flows. **Sovereign AI extends that complexity outward** across regions, providers and regulatory regimes. Running AI in these environments means stitching together infrastructure, cloud, data platforms, models and governance, often across multiple partners. **More than half of organizations (51%)** cite integration complexity in hybrid environments as a top challenge — and it's the number-one challenge overall.

This creates a less obvious dynamic: **The more organizations push for control, the more complex and interdependent their AI ecosystems become.** Progressive organizations are not doing it alone. They are building ecosystems, but with much tighter orchestration, clearer accountability and stronger alignment with their partners.

It is less about owning every component and more about **making a multiprovider system work reliably within defined boundaries.**

Together, private and sovereign AI are changing how AI systems are built, governed and scaled.

Some organizations are already redesigning for that reality. Others will struggle to scale until they do.

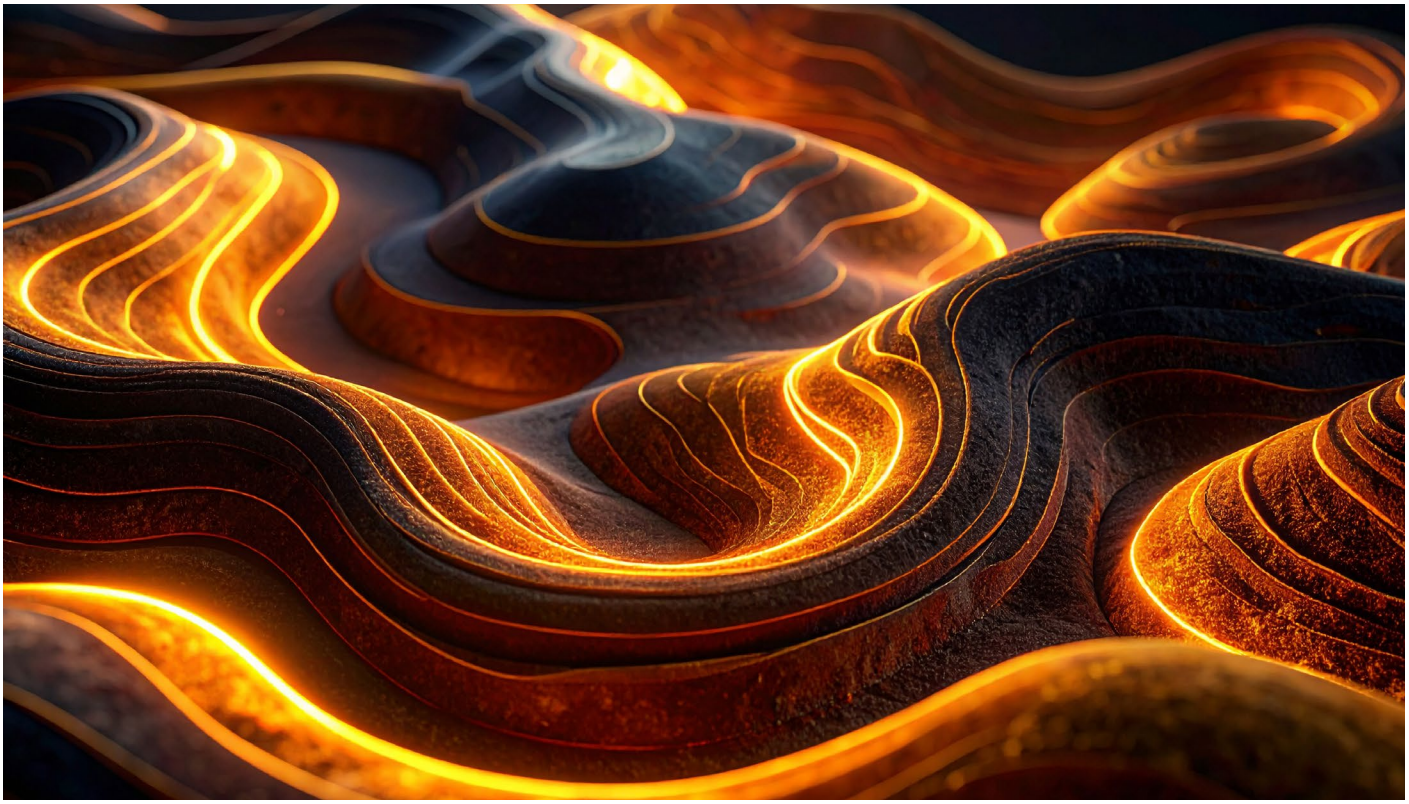

AI is running into a wall — and it's not the model

Enterprise AI strategy has, until now, focused on scale, speed and model performance. The dominant assumptions were simple: Data could move freely, infrastructure could be global and intelligence could be centralized.

Those assumptions are now rapidly breaking down. Today's AI strategies are being tested by deeper infrastructure issues that are as much geopolitical and regulatory as they are technical: where data is stored, who controls the intelligence layer, what happens when jurisdictions diverge and how organizations protect sensitive intellectual property (IP) as geopolitical tensions spill into the technology supply chain.

Think of a multinational bank that operates AI models across regions and suddenly faces conflicting data localization laws and cross-border transfer restrictions. Or a pharmaceutical company training GenAI models on proprietary research that sees their primary concern shift from latency or computing costs to the possible leakage of sensitive trial data or breakthrough formulations.

In both examples, the issue is the same: AI can no longer be separated from questions of control, jurisdiction and exposure. In this context, cross-border data movement becomes a strategic vulnerability. Organizations require private environments that protect sensitive data not only through policy but also through enforceable infrastructure controls.

This NTT DATA playbook, part of our 2026 Global AI Report series, explores how these emerging realities are elevating private and sovereign AI from niche architectural considerations — or simple administrative constraints — to core strategic decisions.

Private and sovereign AI at the heart of strategy

Our research shows that an overwhelming **95%** of organizations now view private or sovereign AI as important to their AI strategy, and **96%** agree — **45%** strongly — that they are considering relocating AI infrastructure to specific geographies because of geopolitical pressures and point-of-origin supply chain concerns.

These considerations generally fall into three broad categories:

1. **Mandated AI sovereignty,** where legal or geopolitical constraints require domestic control
2. **Regulated privacy,** where organizations must demonstrate auditable control over data, models and operations
3. **Strategic AI autonomy,** where global enterprises want greater control over IP, cost and vendor dependence

In the C-suite, the message is even clearer: **98%** say it is imperative to have a private domain that safeguards IP and sensitive data through a publicly nontrainable GenAI model.

95%

of organizations consider private or sovereign AI to be important to their AI strategy.

96%

of organizations agree (**45%** strongly) that they are considering relocating AI infrastructure to specific geographies because of geopolitical concerns.

98%

of the C-suite agree that a private domain that safeguards proprietary IP and data through a GenAI model that is nontrainable publicly is an imperative.

Infrastructure moves into the AI spotlight

In an environment that is more controlled and increasingly localized, what separates organizations that make AI work from those that struggle?

Building private and sovereign AI is more complex than moving workloads to a different location. It requires rethinking the entire stack — infrastructure, data and models — while navigating fragmented rules that vary by region and industry. It also calls for specialized skills, contractual restructuring and, in some cases, new providers. This creates urgency and hesitation at the same time: Organizations understand the risk of inaction but also the complexity of moving forward.

Within the stack, infrastructure has become the real bottleneck. Scaling AI is no longer held back by models but by the environments they depend on. This is where a clear divide is emerging. AI leaders are **21%** more likely than other organizations to be a sovereign approach to how they design and run AI over the next two years, and they build privacy and control into their infrastructure from the start. They also tightly align AI strategy with infrastructure decisions and modernize their environments in step with their AI investment.

AI factories explained

AI factories are purpose-built, high-performance environments designed to continuously train, refine and deploy AI models at scale. They integrate specialized hardware and software and often orchestrate AI agents to convert raw data into actionable intelligence. They are relevant to private and sovereign AI because they provide dedicated, automated end-to-end systems capable of operating at the performance and control levels required.

Why infrastructure now determines control

Control is fundamentally an infrastructure issue. Architecture now matters as much as algorithms when it comes to enterprise AI. Once an organization decides that data must remain within defined borders, or that models must operate under stricter governance, the question quickly becomes where and how AI workloads should run.

This is particularly true for cross-border data movement, where architectural decisions directly determine exposure and compliance risk. It places a new requirement on scaling decisions: Before expanding their AI deployments, organizations must demonstrate — both architecturally and conceptually — that sovereignty and governance controls are enforceable.

Legacy infrastructure was built for application transactions and batch processing. AI, by contrast, depends on real-time inference, high-speed data movement, edge processing and tighter control over data, models and execution environments. Imposing these requirements on legacy architectures places strain on outdated systems.

While nearly every organization (**>99%**) is actively reviewing how to integrate AI into legacy environments, **96%** say their existing infrastructure is slowing AI adoption. Furthermore, only **49%** fully agree that their current data infrastructure can support scalable agentic AI deployments.

96%

of organizations agree, 45% very strongly, that legacy infrastructure is slowing AI adoption.

>99%

of organizations are actively reviewing options to integrate AI into legacy infrastructure.

23%

Organizations that incorporate a sovereign approach into their AI strategy are 23% more likely than others to be fully confident that their IT infrastructure will meet their AI needs.

#1

Integration complexity in hybrid and multicloud architectures is the number-one concern in running AI workloads in private environments.

Although organizations taking a sovereign approach to their AI strategy are **23%** more likely than others to have full confidence that their IT infrastructure will meet their AI needs, our data is explicit on the need for help. About 1 in 3 CAIOs (**35%**) identify the difficulty of building, integrating and managing complex AI models in private and sovereign environments as their top barrier to AI adoption — ahead of regulatory uncertainty. This often leads to a need for significant changes to infrastructure and how these environments are managed.

Within this broader infrastructure challenge, data is emerging as one of the most immediate constraints.

About 1 in 3

CAIOs (**35%**) identify the difficulty of building, integrating and managing complex AI models in private and sovereign environments as their top barrier to AI adoption.

Data jurisdiction is becoming an architectural constraint

AI systems depend on large, diverse and often widely distributed datasets. They need continuous access to the data, along with the ability to move and recombine it in real time. However, even in cases where data can currently move legally, it's often without the speed, scale or flexibility that modern AI architectures require.

In addition, as sovereignty evolves from a compliance requirement to a source of differentiation, the need for data privacy and governance also moves to the core of AI design, influencing architectural decisions, operating models and the composition of partner ecosystems.

As AI architectures move away from globally centralized models toward more fragmented, regionally bound systems, organizations are increasingly designing for multiple jurisdictions, each with its own data, regulatory and infrastructure requirements and constraints.

This data-driven shift introduces a new set of trade-offs — between performance and compliance, and between efficiency and control. Models may no longer be trained on the full breadth of global data, and infrastructure is beginning to diverge across regions, adding complexity to already intricate enterprise operations. Most organizations have not had to navigate this at scale before. After years spent integrating siloed systems, they now face a different challenge: rethinking how AI infrastructure is designed, deployed and governed to support private and sovereign AI.

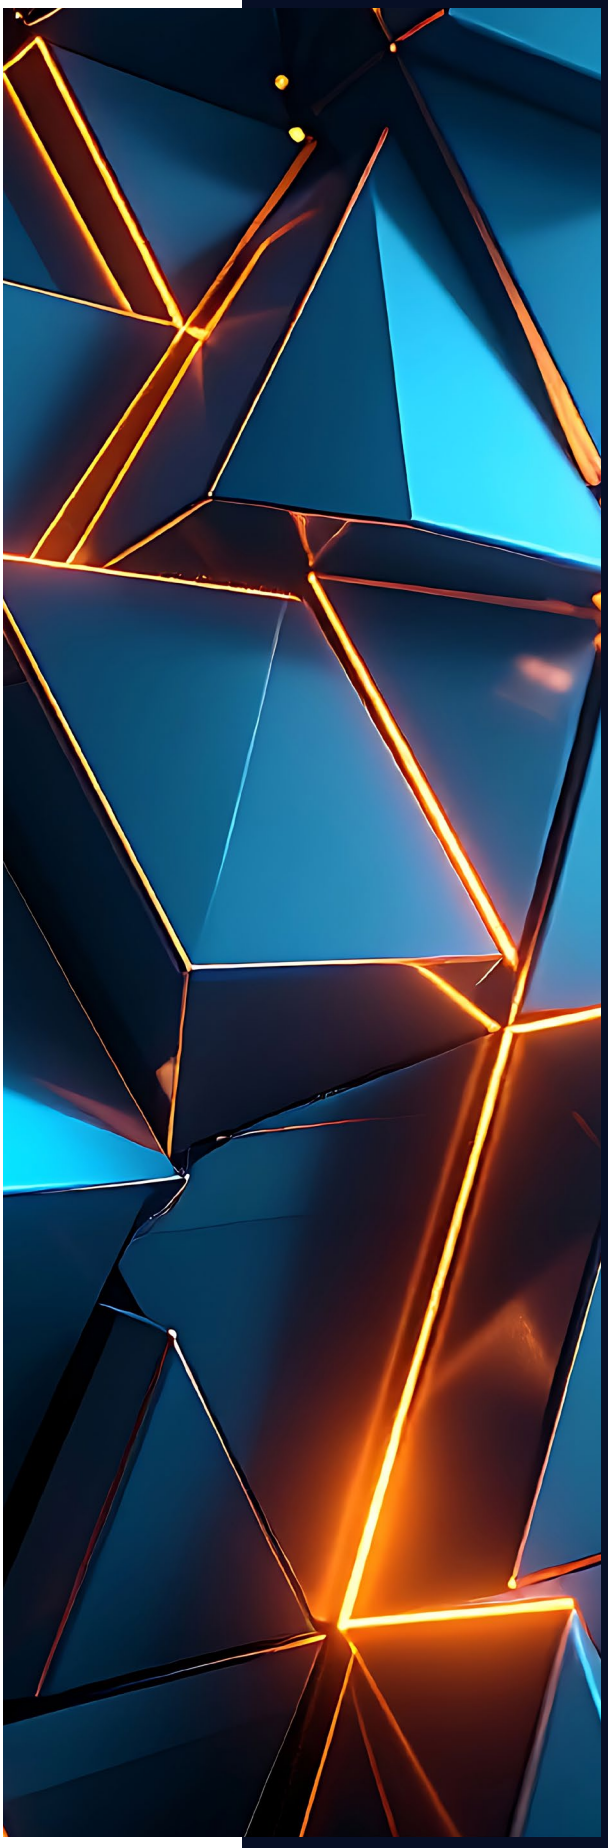

Low confidence in meeting data sovereignty needs

Globally, **96%** of organizations fear privacy violations and misuse of customer data related to the use of AI and GenAI, yet fewer than half (**47%**) have full confidence they can meet data sovereignty requirements.

Keeping data local may sound straightforward, but in practice it requires storage capacity, processing power and resilient networks within defined boundaries. It is not a policy decision alone but also a major capital and operational commitment. Fewer than half of organizations (**48%**) strongly agree they have invested sufficiently in data storage and processing capacity to support AI workloads.

In addition, a majority of CEOs (**57%**) rank data privacy and sovereignty across geographies and environments as a risk for their organization — and it's the number-one security or compliance governance threat overall.

96%

of organizations fear privacy violations and the misuse of customer data from AI and GenAI.

#1

Data security (including data privacy) is ranked as the top challenge in AI adoption, excluding budget concerns.

Only

47%

have full confidence that they can meet data sovereignty needs for AI.

#1

Ensuring data privacy and sovereignty across geographies and clouds is the number-one security or compliance governance threat affecting organizations.

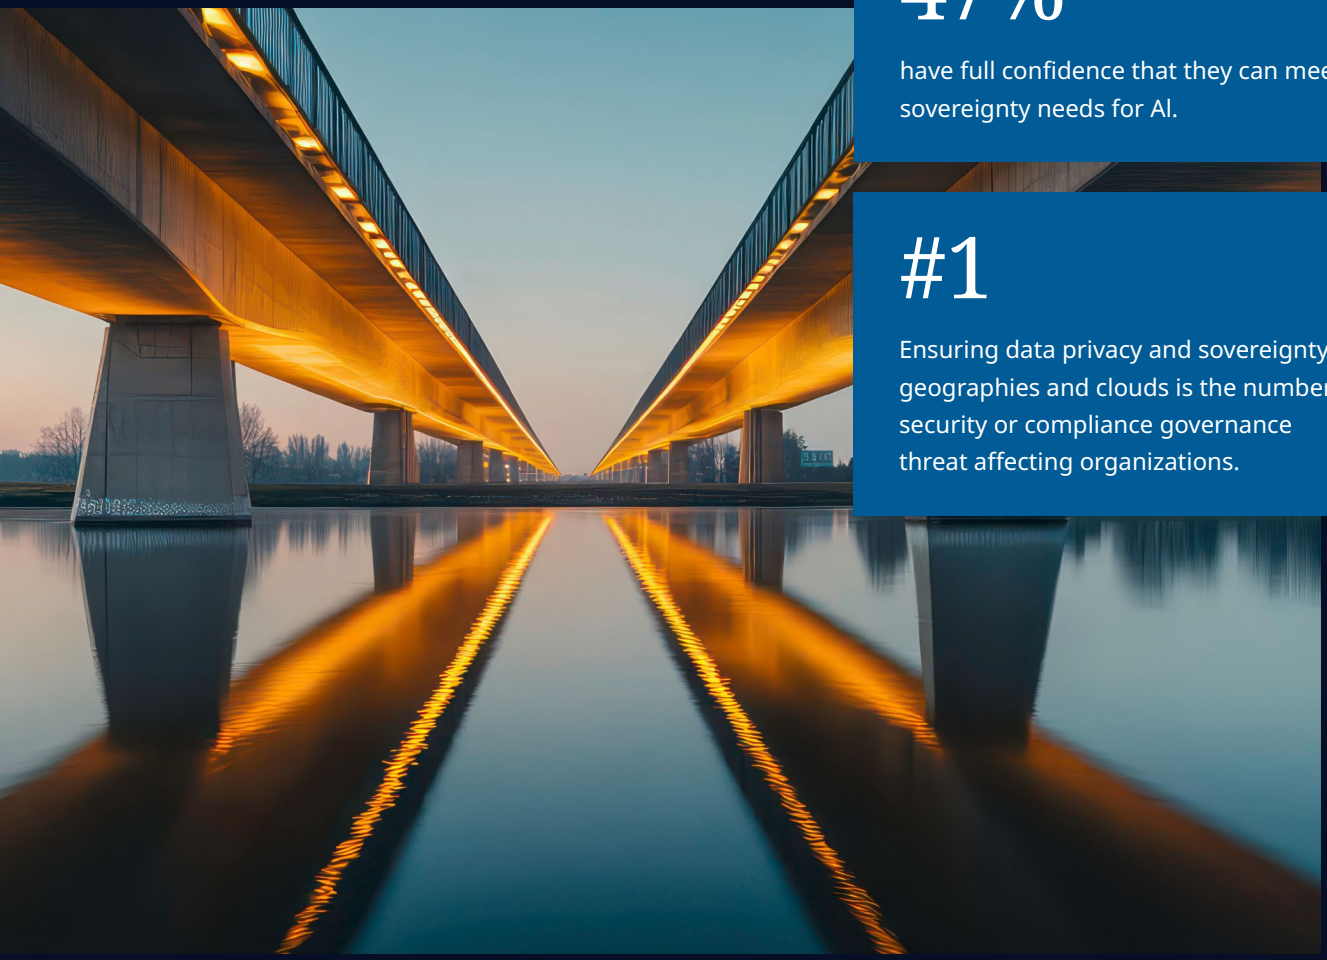

A layered approach to sovereign AI

At the core of sovereignty is deliberate design across three layers: infrastructure, data and models. Infrastructure sovereignty concerns who controls the computing and the platforms. Data sovereignty focuses on who controls the data, where it resides and under which legal framework it is processed. Model sovereignty addresses who controls how intelligence is trained, tuned and distributed. True sovereignty means keeping all three layers local.

Data jurisdiction sits at the center of this stack. The infrastructure, data and model layers are interconnected, but data is often the limiting factor: It determines what must stay local, what can be shared and where intelligence can operate. An organization might keep their data within national borders and retain tight control over their infrastructure, but if the model is governed elsewhere, control over the intelligence layer remains incomplete.

This is why many organizations are starting to deliberately separate data from intelligence. They keep data local within geographic or regulatory boundaries while allowing models, algorithms or learned capabilities to be shared more broadly — through regional fine-tuning, localized data pipelines, federated learning or strict residency controls combined with globally orchestrated AI services.

Data security must be engineered into architecture

Jurisdictional control is adding a new dynamic to data security. Large AI models depend on extensive data pipelines, third-party integrations and continuous retraining cycles, each of which introduces the possibility of leakage, misuse or simple misconfiguration.

Addressing these risks demands unified data governance: deliberate data classification, segmented access using zero trust principles, encryption across the data lifecycle, lineage tracking, regular red teaming, and centralized identity and key management.

In private and sovereign AI, control is meaningful only if it can be enforced through how the environment is built and operated.

Much to consider in workload decisions

All of these realities determine organizations' AI infrastructure choices. Data jurisdiction is no longer a downstream compliance requirement. Instead, it is becoming an upstream design constraint.

Different AI workloads impose different demands on computing density, network design, resilience and data gravity, and compliance can make where a workload runs just as important as how it runs. Many organizations are therefore adopting hybrid architectures, reserving controlled environments for sensitive data, deterministic performance and regulatory oversight while using other environments for data where the risk is lower.

Our data shows that **97%** of organizations agree that critical workloads are best kept in private or on-premises environments, with less sensitive or noncore tasks handled elsewhere. Cost predictability is another significant factor in these decisions.

The conclusion is straightforward: Organizations that don't control their data and infrastructure don't control their AI.

AI leaders are already redesigning their infrastructure accordingly, and our research suggests that their disciplined approach to control is becoming a lever for performance improvement.

Everyone sees the shift, but few are acting on it

There is little doubt that organizations understand the importance of private and sovereign AI. The issue is that many are not acting on it with the required urgency and discipline.

Across the market, momentum still trails intent. Our data shows only **29%** of all organizations are prioritizing sovereign AI in a concrete, near-term way.

It's a familiar challenge: While most organizations recognize the need to run AI in controlled and compliant environments, many are still working out how to turn that into scalable architecture and operating models, especially as private and sovereign requirements add complexity. The shortfall is both in execution and in confidence.

The action gap also varies by region. In the European Union, sovereign AI investment is often more explicitly driven by regulation, while in parts of the Middle East, political and national strategy considerations may outweigh regulatory concerns. This means organizations are not responding to a single global mandate but to a patchwork of local expectations, sector pressures and trust requirements.

Furthermore, industry and use case frequently matter as much as national borders. Organizations that align their infrastructure choices with local expectations can strengthen customer trust and public-sector credibility. Acting on sovereign AI means more than declaring intent. It requires visible choices — such as demonstrable data residency, in-country partnerships and alignment with national digital strategies — that turn sovereignty from a stated priority into an operating model.

Industry and use-case factors influencing AI infrastructure design choices

Geography sets the backdrop, but industry risk profiles and use cases often determine how strongly sovereignty shapes AI architecture.

Our research shows that some industries are far more likely to consider a sovereign approach to designing and deploying AI solutions in the next two years. Globally, public-sector organizations lead at **37%**, followed closely by healthcare (**35%**), natural resources (mining, oil and gas) (**35%**) and manufacturing (**33%**).

In these industries, the consequences of failure are clear. In the public sector, AI systems may connect to national security, citizen data or critical public services. A breach or disruption can undermine public trust or compromise state resilience.

And in healthcare, AI increasingly supports diagnostics, treatment planning and clinical research. The idea of sensitive medical data flowing freely across jurisdictions raises immediate ethical and regulatory red flags.

Industries likely to consider a sovereign approach to designing and deploying AI solutions in the next two years (global)

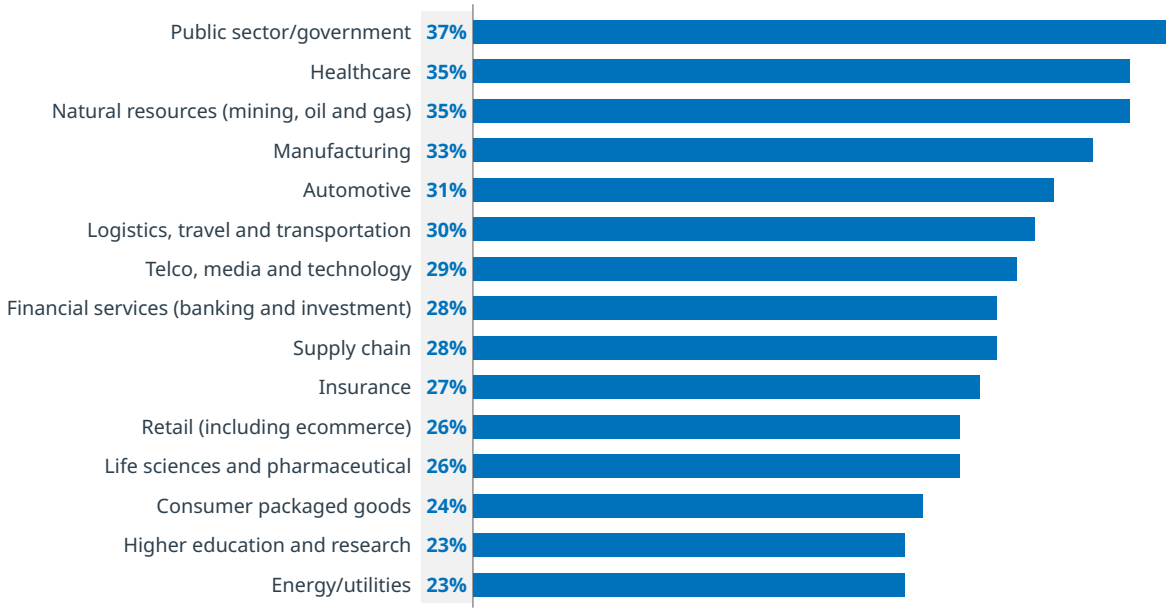

Which approaches are your organization most likely to follow for designing and deploying AI solutions over the next two years?
Base: All respondents, excluding "Don't know" responses (n=2,567)

Regional differences

At a regional level — and, often more explicitly, at a country level — there is some variation in the ranking of industries that are likely to consider a sovereign AI approach. For example, manufacturing moves into the number-one spot in Asia Pacific, the public sector and government are in the lead in Europe, and healthcare is ranked first in North America — a sign that local or regional factors also affect strategic AI deployment choices.

Cloud security posture is lagging, too

AI ambition is also accelerating faster than security readiness. Just **38%** of respondents say they feel very confident in their current cloud security posture, and fewer than half (**48%**) report being highly prepared to manage AI and cloud security risks, with formal plans in place.

Many organizations are moving quickly on AI while still fortifying the guardrails that private and sovereign environments demand. This makes the transition more difficult and increases the premium on partners that can bring implementation discipline as well as technical depth.

Cloud security is foundational to private and sovereign AI because it enforces data residency, access control and compliance requirements while protecting sensitive models and datasets from unauthorized access. Without a strong cloud security posture, organizations risk exposing their AI systems to breaches and governance failures that undermine the very sovereignty and privacy they aim to achieve.

Just

38%

of organizations are very confident in their cloud security posture.

And only

48%

say they're highly prepared to manage cloud and AI security risks with formal risk management plans.

Leaders redesign early and move decisively, creating competitive divergence

AI leaders place private and sovereign AI at the core of their operating models, treating it as a foundational design principle. They see it as a driver of competitive advantage, greater negotiating power and deeper in-house expertise, and they're aligning their overall business strategies accordingly.

Their pace is starting to pull them ahead. A gap is opening up, both structurally and technologically, between these leaders and organizations that are taking a more incremental path as they try to stretch existing setups to fit private and sovereign demands. Leaders are consistently ahead — in some cases, by 10 to 11 percentage points — across prioritization and readiness indicators.

Percentage of organizations that have fully aligned their IT infrastructure strategy with their AI strategy

How aligned is your organization's AI strategy with your IT infrastructure strategy?
Base: All respondents, excluding "Don't know" responses (n=2,526)

Percentage of organizations that say sovereign and/or private AI is an extremely important factor in their AI strategy

How important are the following factors in your AI strategy? – Sovereign/private AI.
Base: All respondents, excluding "Don't know" responses (n=2,567)

Percentage of organizations that see sovereign AI driving a competitive advantage, negotiating power and expertise

What is the key driver to your organization investing in a sovereign AI solution?
Base: All respondents, excluding "Don't know" responses (n=2,567)

A new approach to AI governance

As AI leaders turn the spotlight on infrastructure, they are also taking a more structured approach to AI governance, which they embed from the outset of their AI initiatives.

Although confidence in governance remains uneven across the market, these leaders are far more likely to pursue centralized governance while recognizing the need for federated operating models. Data privacy and sovereignty rank among their top governance concerns, and they are more likely than other organizations to formalize accountability through executive-backed steering committees that bring together business, legal and security stakeholders.

Percentage of organizations that already follow a centralized AI governance model

Which approach best describes your AI governance model?
Base: All respondents, excluding "Don't know" responses (n=2,567)

Percentage of organizations that say cross-geography data privacy and sovereignty is a top governance concern posing a threat to their organization

When scaling AI across your organization, which of the following security or compliance risks pose the greatest threats to your organization?
Base: All respondents, excluding "Don't know" responses (n=2,567)

Percentage of organizations that have an AI steering committee in place with an executive sponsor, domain leaders, and representation from legal, security and more

Which of the following have been initiated as part of your AI governance model?
Base: All respondents with an AI governance model, excluding "Don't know" responses (n=2,505)

The CIO's view of sovereign AI

For CIOs, a sovereign AI strategy can be both strategic and defensive. According to the CIOs surveyed globally for our report, gaining a competitive advantage, bargaining power and in-house expertise are the main drivers for investing in sovereign AI, along with meeting in-country regulatory and data requirements, the desire to simplify the IT landscape and reduce vulnerability, and the need for national security and strategic autonomy.

This differs from the organization-wide view, where simplifying the IT ecosystem to reduce risk is ranked as the second-most important driver of investment in sovereign AI.

CIOs' top drivers for investing in a sovereign AI solution

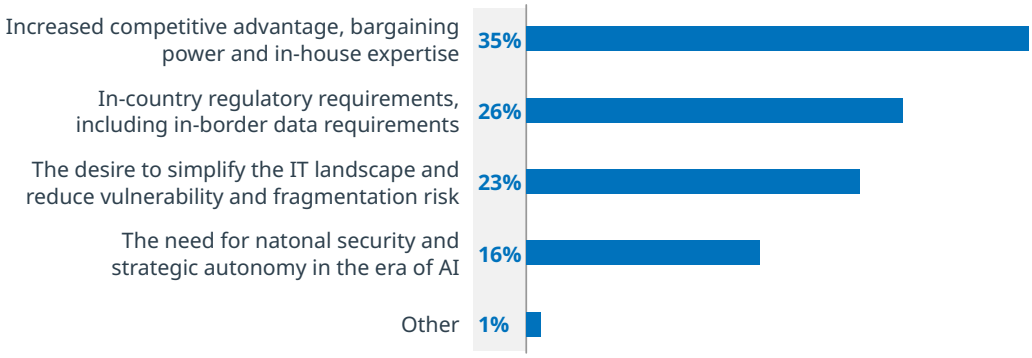

What is the key driver to your organization investing in a sovereign AI solution?
Base: All CIO respondents, excluding "Don't know" responses (n = 211)

Private and sovereign AI sound like independence but are built on tightly orchestrated ecosystems

Although private and sovereign AI are becoming strategic necessities, they are not becoming simpler. **Greater control often increases complexity as AI ecosystems become more interconnected and interdependent.**

Designing and running private and sovereign AI environments requires coordination across multiple providers, infrastructure, data, platforms and models while accounting for governance, security and compliance within defined boundaries. Until now, most organizations have had little reason to master this level of integration, as our research shows:

- Among those advancing sovereign AI strategies for GenAI, **40%** of organizations cite infrastructure modernization as their single biggest challenge, followed closely by the difficulty of assessing and integrating complementary technologies.
- More than half of organizations (**51%**) cite integration complexity in hybrid environments as a top challenge when running AI workloads on private cloud, ranking it as the number-one challenge overall.
- And only **49%** say they are fully satisfied with the data governance and sovereignty capabilities of their existing cloud environments. Those who are not fully satisfied point first to a lack of cross-company collaboration, then to management challenges and architectural complexity, and then to rising costs as the reasons for their discontent.

In addition, privacy and sovereignty cut across legal, IT, risk, procurement and business units. The operating model challenge is therefore organizational as well as technical.

These implementation issues sit at the heart of whether private and sovereign AI can move from aspiration to execution.

Developing the partner ecosystem

Private and sovereign AI cannot be executed as a single-organization capability. Because these capabilities take time to build, organizations that establish partner ecosystems early are better positioned to scale without disruption.

Our research shows that **32%** of CAIOs — and **38%** of AI leaders — consider the availability of a private or sovereign AI solution as a key factor when choosing enterprise AI platforms. But selecting a platform is only the beginning. Implementing it effectively creates a new layer of complexity, with governance, orchestration, lifecycle management and integration with existing systems all coming into play.

About 1 in 3

CAIOs (**32%**) say that the availability of a private or sovereign AI solution is a top factor when considering an enterprise AI solution.

Dealing with this complexity increasingly depends on collaboration with partners that operate across the full stack — physical data centers, AI infrastructure, platforms, models, security and higher-level AI services — while also connecting those capabilities to sector-specific requirements.

The value of these partners lies in their ability to both provide the technology provision and translate policy and control requirements into production-ready environments that can run at scale. They orchestrate complex multivendor ecosystems instead of optimizing only a single component, and they remain engaged throughout the lifecycle — from strategy and design to build, deployment, operations and ongoing optimization. Experienced ecosystem partners, proven reference architectures, deep provider relationships and repeatable modernization playbooks shorten timelines and limit disruption.

This is what ecosystem orchestration now means in AI: not simply assembling technologies but aligning architecture, governance, security and operations across providers, regions and regulatory contexts.

> “Private and sovereign AI cannot be executed as a single-organization capability.”

What effective partners do

- Operate across the full stack, from infrastructure to models and industry workflows.
- Support the full lifecycle, from strategy and design to build, run and optimization.
- Orchestrate multivendor ecosystems across providers, regions and regulatory contexts.
- Combine global scale with local execution and jurisdiction-aware controls.

AI leaders reflect this shift in how they build partner ecosystems. They are more likely to factor private and sovereign AI capabilities into vendor selection, demand more flexibility in contracts and adapt more quickly to changes in their relationships with large-scale providers. More mature organizations place greater weight on partners with proven sovereign capabilities, particularly in data governance.

Percentage of organizations that view the availability of a private or sovereign AI solution as an important factor in their decision-making process when considering an enterprise AI solution

When considering an enterprise AI solution, what are the important factors in your decision-making process?
Base: All respondents, excluding “Don’t know” responses (n=2,566)

In short, as the market moves from AI experimentation to AI operating models, partner strategy becomes a core part of AI strategy. The organizations that scale successfully will be those that combine control with orchestrated ecosystems, turning complexity into trusted execution. It is less about owning every component and more about **making a multiprovider system work reliably within defined boundaries.**

Start the next chapter of AI

AI is entering a new phase that is defined less by model innovation and more by the conditions under which those models can operate.

The constraints are becoming harder to ignore. Infrastructure is under strain, data can no longer move as freely, and sovereignty and privacy requirements are starting to influence how AI systems are built and deployed. What once felt like a given — global scale, centralized intelligence and seamless data flows — now has to be deliberately designed, enforced and proven.

AI leaders are already adjusting. They are aligning their infrastructure and AI strategies and redesigning their architectures to build in control, locality and governance from the start. Other organizations are moving more cautiously. However, as privacy and sovereignty demands tighten, this gap is becoming harder to close.

The implication is clear: Scaling AI is no longer just a question of capability but also one of design. Organizations need to decide where AI runs, how data is handled and what can be controlled, across infrastructure, models and ecosystems.

Access alone won't define this next chapter. How intentionally AI is built will matter just as much, and early movers will be better placed to scale with confidence in a world where privacy and sovereignty matter more than ever.

Visit our website to see how NTT DATA can help you chart a path forward with AI.

Explore our research data in detail

Our 2026 Global AI Report is another milestone of primary research and thought-leadership from NTT DATA. Additional expert and executive insights will follow. Ask us how our comprehensive global research data, coupled with our consulting and services expertise, can support your organization's success, and look out for more research-driven thought leadership soon.

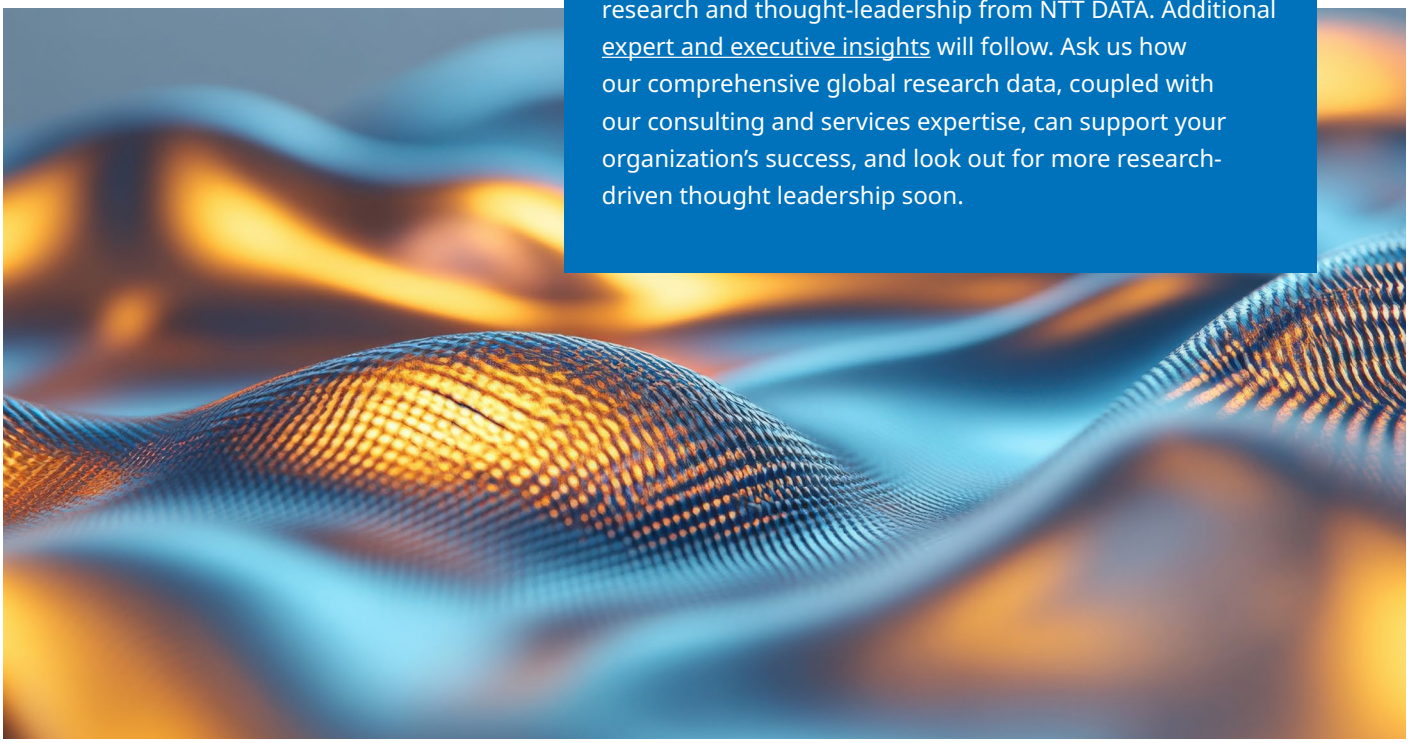

Definitions

Private and sovereign AI

Private and sovereign AI refer to requirements for countries or organizations to independently develop and deploy AI. Depending on specific needs for autonomy or regulatory compliance, or legal and geopolitical constraints, increasing levels of privacy and sovereignty are applied.

4 key dimensions of private and sovereign AI

01
Territorial
Residency of data and computing

02
Operational
Management, control and security of systems

03
Technological
Ownership of code, models and IP

04
Legal and regulatory
Jurisdiction over data access, use and rights

Private AI puts sensitive IP and regulated data in controlled, dedicated AI environments with strict identity and access management, encrypted data pipelines and model lifecycle controls.

Sovereign AI keeps data, infrastructure, computing and control within national or regional borders, often in response to geopolitical realities.

AI maturity

In NTT DATA's global AI research, AI leaders and AI laggards (the opposite cohort) are found in every revenue range and in every industry. What separates leaders from laggards is a defined AI strategy and a higher level of AI maturity.

Levels of AI maturity defined

- **No plans:** Use of AI not yet explored in the organization
- **Explorer:** Strategies and plans under consideration, but no adoption or capability so far
- **Novice:** Just starting, with limited experience and use cases
- **Enabled:** Sporadic and somewhat siloed use, with feasibility pilots and limited adoption by individual business units in mostly noncore functions
- **Mature:** Broad and strategic use across business units and functions, with strong governance, best practices and scalable workloads
- **Evolved:** Incorporated into core and noncore business functions as well as continuous service delivery; AI-led innovation accelerates business transformation and advances business outcomes

Read more: The **2026 Global AI Report: A Playbook for AI Leaders**, the first publication in this series, presents nine characteristics of AI leaders to illustrate how they operate and why they are already seeing the benefits of their investment in AI.

About the research

Research methodology

All content in the playbook is based on independently sourced research data.

Primary research fieldwork was conducted for NTT DATA by STRAT7 Jigsaw, an international strategic-insight, analytics and market intelligence agency with an exclusively senior team.

Participants were prescreened and then selected via random sampling on the basis that they had influence or decision-making authority on their organization's AI and/or technology strategy. The research findings are derived from two distinct data sets, gathered via online questionnaires that ran in September and October 2025.

Data integrity, validation and analysis were performed by NTT DATA's specialist in-house Primary Research and Benchmarking Team in conjunction with STRAT7 Jigsaw. Data and outliers were validated in conjunction with STRAT7 Jigsaw and in accordance with standard research-industry rules, disciplines and best-practice approaches. The data is presented at a 99% confidence level with a 3% margin of error.

Globally, the research spans more than 30 markets in five regions, across more than a dozen industries and a combined sample of nearly 5,000 senior decision-makers and influencers. The samples are weighted heavily toward large organizations and senior decision-makers, with strong C-suite representation.

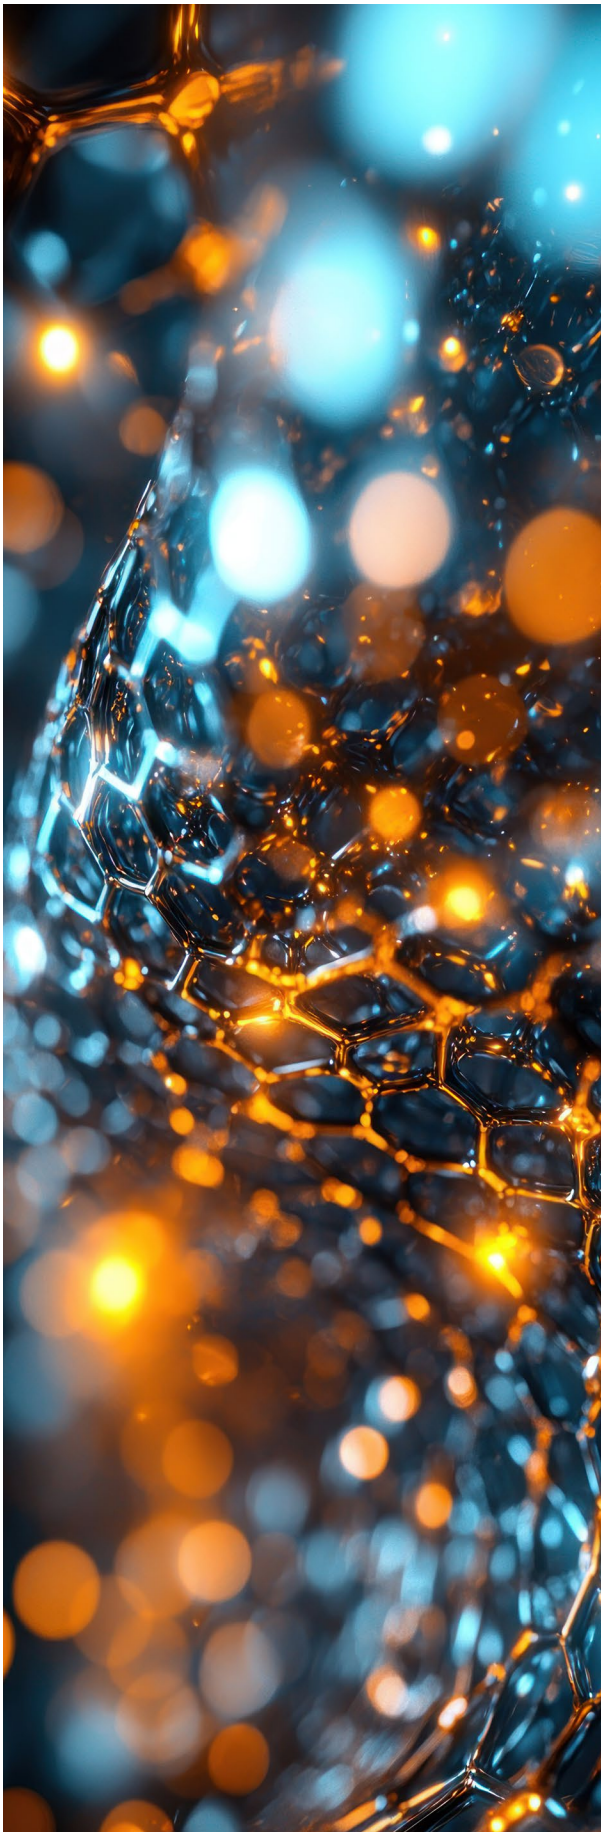

Our research in numbers

2025 global AI research

- A balanced sample of 2,567 AI decision-makers (94%) and influencers (6%).
- Coverage spans 35 markets in 5 regions and 15 industries.
- 84% of respondents were from large organizations with more than 10,000 employees; 11% from organizations with 5,001 to 10,000 employees; and 5% from organizations with 2,500 to 5,000 employees.
- 79% of respondents were from the C-suite; 15% were at Senior Vice President, Vice President, Head of or Director level; and 6% were Senior Managers or Specialists.
- 10% of respondents were CEOs and 4% were CAIOs. Additionally, 31% were from IT business functions (excluding security); 11% from IT security; 6% from digital; 21% from operations; and 17% from business support (including legal, compliance, risk, finance, marketing and HR functions).

2025 technology architecture research

- A balanced sample of 2,335 technology decision-makers (94%) and influencers (6%).
- Coverage spans 33 markets in 5 regions and 13 industries.
- 83% of respondents were from large organizations with more than 10,000 employees; 11% from organizations with 5,001 to 10,000 employees; and 6% from organizations with 2,500 to 5,000 employees.
- 73% of respondents were from the C-suite; 21% were at Senior Vice President, Vice President, Head of or Director level; and 6% were Senior Managers or Specialists.
- 12% of respondents were CEOs, and 5% were CAIOs. Additionally, 46% were from IT business functions (excluding security); 9% from IT security; 7% from digital; 12% from operations; and 10% from business support (including legal, compliance, risk, finance and HR functions).[2]

[2] Percentages are rounded to whole numbers; as a result, totals may not add up to exactly 100%.

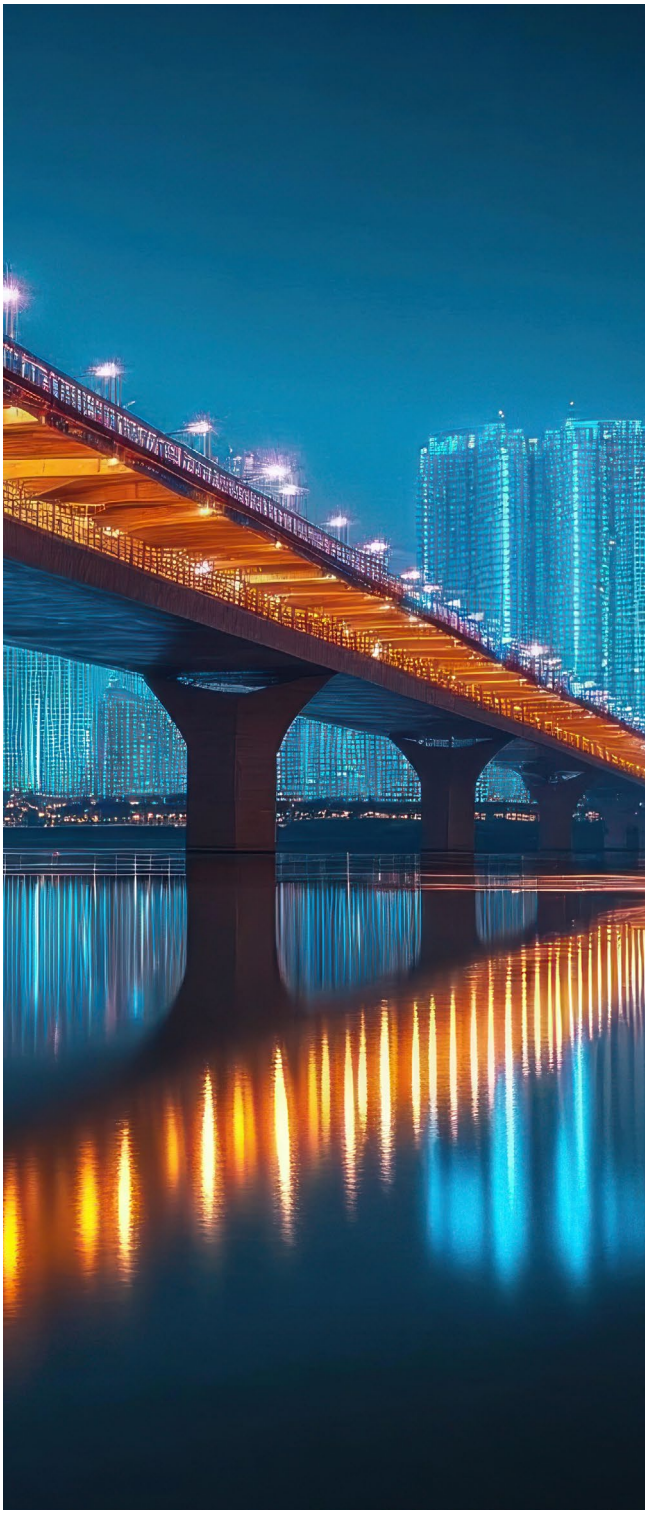

Redesign and scale for the new AI reality

NTT DATA is trusted by 75% of the Fortune Global 100 to drive transformation and growth. We empower organizations to innovate faster and operate smarter through an end-to-end portfolio spanning consulting, applications, data and AI, digital workplace and secure digital infrastructure. As demand for private and sovereign AI accelerates, we deliver trusted services and infrastructure that enable data control, regulatory alignment, and confident AI at scale. NTT DATA is part of the $90+ billion NTT Group, which invests over $3 billion annually in R&D to help organizations move boldly and sustainably into the digital future.


NTT DATA