MIT Sloan Management Review

What CEOs Need to Know About Sovereign AI

Mauro Macchi, Ajoy Menon, Mauro Capo, and Surya Mukherjee

Chief executives can balance the local ambitions of countries with the global scale of innovation in a changing geopolitical environment.

July 16, 2026
Reprint 68116

sloanreview.mit.edu

What CEOs Need to Know About Sovereign AI

Chief executives can balance the local ambitions of countries with the global scale of innovation in a changing geopolitical environment.

By Mauro Macchi, Ajoy Menon, Mauro Capo, and Surya Mukherjee

As multinational companies implement artificial intelligence workflows and look to adopt AI across their global operations, they are increasingly running up against country-specific regulations and policies that aim to govern AI and align its use with national priorities and local cultural norms. These regulations and policies — which fall under the umbrella of "sovereign AI" — govern where data is stored and processed, whose infrastructure is used for training and operating AI models, and how algorithmic decisions are reviewed and enforced in a given jurisdiction. Many markets are now developing their own sovereign AI frameworks to reduce dependence on the United States and China, where nearly 70% of leading AI models originated.

This creates a strategic dilemma for multinationals. Relying on global AI platforms maintains operational consistency but deepens exposure to geopolitical disruption and local market access risk. Localizing data, infrastructure, and models earns regulatory trust but incurs significant cost and complexity when a company operates across dozens of jurisdictions with differing requirements. The challenge is that policies vary significantly by country and are evolving rapidly, making a single global AI strategy untenable, and fully independent local systems impractical.

Most companies are responding defensively, treating sovereign AI as a compliance obligation managed by legal or IT teams. Our December 2025 survey of 1,928 executives across 28 countries reveals a striking gap: Sixty percent of respondents said that rising geopolitical risk makes them more likely to pursue sovereign technology solutions, yet only 15% have made AI sovereignty a CEO or board-level priority, and fewer than 13% see it as a growth driver rather than a cost.

In this article, we argue that sovereignty is better understood as a continuum of choices and that the companies best positioned to scale AI globally are those that treat those choices as a source of competitive advantage rather than a constraint to be minimized.

The Sovereign AI Landscape

The regulatory landscape governing AI has shifted significantly in recent years, from relatively narrow data residency rules to a far broader set of requirements governing models, infrastructure, and how algorithmic decisions are made and enforced. Most major markets are now developing their own sovereign AI frameworks, resulting in a patchwork of locally governed ecosystems, each with distinct rules, data standards, and expectations for responsible use.

The complexity is already visible at every layer of the technology stack. Companies operating in European Union member states must conduct formal risk assessments, maintain detailed technical documentation, and submit it to national market surveillance authorities under the EU AI Act. (The AI Act is now in active enforcement, with its most comprehensive requirements for high-risk AI systems taking effect in August 2026.) Simultaneously, they must comply with prior standards and policies, like GDPR (General Data Protection Regulation), NIS2 (Network and Information Security 2), and DORA (Digital Operational Resilience Act). Companies must also prepare to align with the sovereignty package announced by the European Commission in June 2026, which includes two legislative proposals and a strategic road map to bolster the EU's AI sovereignty.

Meanwhile, companies looking to do business in Saudi Arabia must navigate strict data localization requirements — including obligations to store nationally sensitive data within the country — alongside cross-border transfer rules that require adequacy assessments

Sovereign AI is a strategic bet involving geopolitics, capital allocation, supply chain resilience, and long-term competitiveness.

or contractual safeguards, all within a governance framework that is still taking shape. There is no dedicated AI law, and binding obligations currently flow from data protection and cybersecurity regulations rather than AI-specific legislation. The same governments driving these requirements are also pouring billions of dollars into building the infrastructure and incentives to enable sovereign AI solutions.

Companies have started to develop strategies to navigate this fragmented landscape. We know from our consulting work that three global banks are rethinking their tech strategy in the EU: They're limiting further migration of sensitive systems into foreign public cloud systems and instead building a shared platform in their home countries.

In the U.K., senior leaders of multinational banks are exploring a domestic alternative to Visa and Mastercard to reduce reliance on U.S.-owned payment networks. These are early signals of a structural shift in how multinationals must think about AI infrastructure. This raises an urgent question for CEOs: How do you scale AI globally when the rules governing it are local, fragmented, and still being written?

To answer this question, CEOs need to make three strategic choices: where accountability for decisions should sit, how much sovereignty their operations require, and which external partners can help them execute.

1. Make sovereignty a strategic priority.

The first and most urgent CEO decision is raising AI sovereignty to the level of a strategic concern. Our survey found that most organizations delegate decisions regarding AI sovereignty to chief data/AI officers (37%) or compliance/risk officers (28%), while only 15% of organizations have made it a CEO- or board-level priority. When sovereignty sits in IT or compliance, it results in fragmented decisions across business units, inconsistent approaches across markets, and missed opportunities to turn sovereignty into local advantage. Sovereign AI is not an IT architecture choice. It is a strategic bet involving geopolitics, capital allocation, supply chain resilience, and long-term competitiveness, and the decisions it requires can be made only at the top.

What does that look like in practice? Consider BNP Paribas, one of the largest banks in the EU. Since 2023, it has built a deepening partnership with Mistral AI, Europe's leading sovereign AI model provider, culminating in a groupwide multiyear agreement in 2024 and a renewed three-year extension in 2025 covering software, co-development research, and on-premises deployment. The bank also backed Mistral financially, participating in both its 385 million euro ($445 million) funding round in 2023 and its $640 million Series B in 2024. The partnership is driven by the C-suite with sovereignty as a key consideration. Keeping AI on-premise, under the bank's direct control, ensures sensitive data stays within European regulatory jurisdiction. Such decisions — say, which AI ecosystem to depend on, which regulatory frameworks to operate within, how to manage geopolitical exposure — carry implications for capital allocation, competitive positioning and long-term resilience. They belong on the C-suite and board agenda.

2. Treat sovereignty as a continuum.

As such CEO-level choices demonstrate, sovereign AI does not require full independence or the wholesale replacement of existing systems. CEOs must recognize that sovereignty is a continuum. Depending on the industry, national context, and specific use case, aspects of the company's technology stack may need varying degrees of adjustment to meet sovereignty concerns. Only about one-third of all executives we surveyed believe that their AI workloads require any such sovereign adjustment. Those actions can involve anything from implementing targeted safeguards around data residency or legal oversight to replacing entire elements of the technology stack, such as models or infrastructure, to satisfy sovereignty requirements.

The degree to which sovereignty matters to the design of a company's AI system rests on three considerations.

The first is industry risk. Defense, health care, energy, and financial services are more sensitive industries because AI systems have implications for national security, citizen safety, and economic stability. Governments are likely to have greater AI sovereignty concerns for companies operating in those domains. In contrast, retail, tourism, and consumer services can often rely on global platforms when strong local safeguards are in place.

The second is national context. Countries pursue different strategies shaped by geopolitics and economic priorities. China has built a full-stack China-for-China model. Singapore emphasizes interoperability and cross-border trust. The U.K. and much of Europe favor hybrid approaches that combine global platforms with selective local control. Companies operating across these markets must design for variation.

The third and most decisive consideration is the use case. Even within an industry, AI that is used in high-stakes decision-making, such as credit decisions, medical diagnostics, or energy grid optimization, carries far greater risk than AI used in contexts like customer service, marketing personalization, or internal productivity. High-stakes use cases call for greater scrutiny around data governance, model transparency, and regulatory exposure.

All three considerations are relevant to AstraZeneca, a global pharmaceutical company. Pharmaceuticals is a sensitive industry, so AI use carries high stakes by default. But AstraZeneca calibrates its sovereign controls deliberately by national context and use case. In China, adverse drug reaction data must be reported to the National Medical Products Administration and remain within the country under strict governance. To address this requirement, AstraZeneca is using Alibaba Cloud's sovereign infrastructure to deploy local AI models, such as the Qwen large language model, for pharmacovigilance in a private, locally controlled environment.

Outside China, AstraZeneca made a different calculation. For R&D and clinical development — work that is sensitive but not subject to the same regulations as pharmacovigilance — it uses the Amazon Web Services (AWS) public cloud to run large-scale AI and machine learning, including multi-agent systems that enable scientists to query complex biomedical data and generate insights in seconds. The same company, facing different sovereign AI pressures in different contexts, arrived at two different infrastructure decisions.

3. Build hybrid sovereign ecosystems.

Few organizations can, or should, build the full AI stack independently. Our survey found that 55% of organizations plan to use a mix of global and local AI providers, reflecting a shift from dependency toward flexibility.

CEOs must adopt a tailored mix of global and local AI providers, depending on use cases, the regulatory landscape, and risk tolerance.

For organizations seeking global scale, partnering with hyperscalers such as AWS, Google, Microsoft, and Oracle is an attractive option. In software and platforms, Oracle's EU Sovereign Cloud and Microsoft's Delos Cloud partnerships are enabling European companies to run sensitive workloads fully under EU law.

When earning local trust is a priority, partnering with country-endorsed national champions may be the right strategic move. Indosat Ooredoo Hutchison — Indonesia's trusted national telco, jointly owned by Qatari and Hong Kong-based investors — is building Indonesia's first sovereign AI cloud with international partners Accenture and Nvidia while ensuring that national data stays onshore to support local startups and government clients.

Companies that want to avoid full dependence on hyperscalers can access specialized capacity from AI-native players like Nebius, CoreWeave, and Lambda. These providers build infrastructure specifically for AI workloads, allowing companies to access high-performance compute more efficiently, often with regionally deployed capacity that meets local regulatory or data requirements. Helical, a Europe-based biotech company, trains its large-scale biological AI models on Nebius's infrastructure. Nebius is headquartered in Amsterdam and operates data centers across Europe, giving it the compute performance it needs while keeping data within regional boundaries.

When the deepest level of sovereign control is required, and where competitive advantage comes from shared capability rather than proprietary infrastructure, federated consortia offer something other models cannot. The OpenBind Consortium, for example, received 8 million pounds ($10.8 million) from the U.K. government's Sovereign AI Unit to bring together a group of partners that includes the University of Oxford, Diamond Light Source, and MedChemica to build a nationally governed data set for AI-driven drug discovery that is 20 times larger than earlier efforts. The real advantage, however, lies in the sovereign governance model: It allows companies to collaborate, share risk, and scale AI with confidence — something fragmented, ad hoc data sets cannot deliver.

Companies that treat sovereign AI as a strategic design choice rather than a compliance obligation will shape how AI evolves in their markets.

The Strategic Imperative

Most companies focus sovereignty on the data and cloud layers of their tech stack. Our survey found that while 60% apply sovereignty controls to data and 46% to infrastructure, only 22% extend them to AI models themselves, leaving a critical layer exposed. As AI systems become more autonomous and agentic, the decisions and actions that matter most will increasingly happen at the model and agent layers.

This doesn't mean that every company needs full-stack control. Companies need to apply sovereignty to the layers that matter most for their highest-stakes use cases.

In a world of persistent geopolitical uncertainty, sovereign AI has become a competitive capability.

Companies that treat it as a strategic design choice rather than a compliance obligation will shape how AI evolves in their markets. Those that don't risk discovering too late that the intelligence driving their most critical decisions is no longer under their control.

Mauro Macchi is CEO for Europe, Middle East, and Africa (EMEA) at Accenture and chairman of Accenture in Italy. Ajoy Menon is senior managing director and global digital core lead at Accenture. Mauro Capo is managing director and digital sovereignty lead for EMEA at Accenture. Surya Mukherjee is principal director and global sovereign research lead at Accenture Research.

ACKNOWLEDGMENTS

The authors thank David Wood, Bryan Rich, Kunal Shah, and Gargi Chakrabarty for their contributions.

MIT Sloan Management Review

PDFs · Reprints · Permission to Copy · Back Issues

MIT Sloan Management Review articles, permissions, and back issues can be purchased on our website, **shop.sloanreview.mit.edu**.

Reproducing or distributing one or more *MIT Sloan Management Review* articles **requires written permission**.

To request permission, use our website **shop.sloanreview.mit.edu/store/faq** or email **smr-help@mit.edu**