Are You Ready for Sovereign AI?

Governments could function as orchestrator, investor, regulator & anchor customer to drive strategy

By Raju Chellam

The CTO presented the company's AI sovereignty draft plan at the board meeting. "This plan gives us full control of our data, total governance assurance, and independence from external influence," the CTO proclaimed after going through a dozen slides. "I would urge the board to approve it so we can get started."

After a brief bout of bickering, a consensus was reached. "I love it, especially the part where the AI can make autonomous decisions," the chairman said. "Just as long as every decision is reviewed by legal, compliance, risk, audit committee, and especially that one director who only checks email twice a day."

If that anecdote made you snigger, check out this trigger: According to a report by McKinsey, about 30% to 40% of AI spending could be influenced by sovereignty requirements, representing a market of some US$500 billion to US$600 billion globally by 2030. Gartner predicts that 35% of countries will be locked into region-specific AI platforms using proprietary contextual data by 2027, up from just 5% currently.

STRATEGIC SIGNIFICANCE

What exactly is sovereign AI? McKinsey says it does not have a single definition. Rather, it is the result of interactions between four components: territorial (where data and compute physically reside), operational (who manages and secures data and compute), technological (who owns the underlying stack and intellectual property) and legal (which jurisdiction governs access and compliance).

In other words, sovereign AI is about a nation's or an organization's ability to develop and control its own AI capabilities to ensure strategic independence and alignment with domestic values and laws. "Viewed this way, sovereign AI is a spectrum of potential solutions distributed across different tiers of sovereignty, depending on stakeholder and local circumstances," McKinsey explains. "Sovereign AI thus represents one of the largest opportunities within AI."

An effective sovereign ecosystem is not necessarily one in which everything is built domestically. Instead, it is one in which key control points are sovereign by design, even if other elements

> Sovereign AI is not about building everything at home. It is about owning the control points that matter most.

of the stack may remain open to partnerships, interoperability and competition.

"The most effective ecosystems operationalize 'minimum sufficient sovereignty' with a repeatable decision rule: Classify workloads by the importance of regulatory issues and third-party exposure," McKinsey says. "Assign a sovereignty tier with explicit requirements for data residency, key ownership and access controls."

Even jurisdictions with advanced capabilities are rarely self-sufficient across all layers and often rely on external providers in at least part of the stack, particularly in hardware and advanced compute. That's why building a sovereign AI ecosystem involves coordination across four groups: Governments (to shape trust, rules and demand); providers (to create the underlying technology and platforms), enterprises (to convert infrastructure into real economic value) and investors (to supply the capital and risk tolerance needed to scale solutions).

The concern? Companies are changing alignment due to concerns of overly Western influence. AI sovereignty will therefore lead to reduced collaboration as well as duplication of effort. Gartner predicts that nations establishing a sovereign AI stack will need to spend at least 1% of their GDP on AI infrastructure by 2029.

"Countries with digital sovereignty goals are increasing investment in domestic AI stacks as they look for alternatives to the closed US model, including computing power, data centers, infrastructure and models aligned with local laws, culture and region," says Gaurav Gupta, a Gartner vice president. "Trust and cultural fit are emerging as key criteria. Decision makers are prioritizing AI platforms that align with local values, regulatory frameworks and user expectations over those with the largest training datasets."

Localized models could deliver more contextual value and regional LLMs could outperform global models in applications such as education, legal compliance and public services, especially in non-English languages.

CLOUD COMPONENTS

AI is dependent on data, much of which resides in data centers. Data and cloud are the driving forces powering digital transformation and business model reinvention. The problem? High risk and reliance on third-party service providers, data stored on external platforms, and potentially an ever-larger surface for cyberattack.

"Regulators across Europe, the Middle East and Africa are stepping up their focus on data residency, data protection and other key aspects of cloud sovereignty," says a report by PwC. "But sovereignty is far more than just a compliance exercise. As organizations move beyond cloud migration to focus on optimization, trust and accountability, sovereignty allows them to innovate on their own terms. They can leverage their cloud and AI capabilities, safe in the knowledge that data access is under their control and their tech choices remain open."

Should the formulation of your sovereign cloud strategy be left to tech teams? Nope. PwC advocates active input and sponsorship from business, compliance and data management teams and urges organizations to define (their sovereign strategy), tailor (AI models to strategic goals), set (clear and realistic boundaries), design (their cloud architecture) and build (strategic partnerships with key stakeholders).

"As organizations embrace digital transformation, clear control over where and how data is governed should be a strategic imperative," says James Rashleigh, a PwC UK cybersecurity partner. "By taking ownership of cloud data governance, leadership teams will not only be able to deliver compliance, but also resilience."

Data centers are the critical backbone to help enable AI sovereignty. "As a result, data centers and AI factory infrastructure will see explosive build-up and investment, propelling a few companies that control the AI stack to achieve double-digit, trillion-dollar valuations," Gartner's Gupta says.

So how can your organization be ready for sovereign AI? Gartner advises focusing on the following:

- **Design:** Model agnostic workflows using orchestration layers that enable you to switch between LLMs across regions and vendors.
- **Ensure:** AI governance, data residence and model tuning practices are able to meet country-

> Governments may shape sovereign AI by setting rules, funding infrastructure and creating trusted demand at scale.

specific legal, cultural and linguistic regulations and requirements.
- **Establish:** Relationships with national cloud providers, regional LLM vendors and sovereign AI stack leaders in priority markets, as well as build a vetted list of partners.
- **Monitor:** AI legislation, data sovereignty rules and emerging standards that may affect where and how they can deploy AI models and process users' data.

GOVERNMENT GOALS

Governments could function as orchestrator, investor, regulator and anchor customer. Because they possess the unique ability to turn fragmented ambition into coordinated execution and also set AI sovereignty goalposts.

Governments define which workloads require strong sovereignty (such as defense, sensitive citizen data apps and critical infrastructure), which can use hybrid models, and which can remain largely global. They then translate those choices into actionable controls (such as data classification, key ownership and auditability). By creating certification regimes, governments can help standardize what "trusted" means so that regulated industries can adopt the guidelines quickly.

McKinsey says successful sovereign AI ecosystems tend to emerge through three overlapping waves:
- **The First Wave:** Focuses on establishing the baseline and unlocking early demand. Leaders clarify which workloads require sovereign controls, translate those decisions into governance and procurement mechanisms, and launch a small number of lighthouse use cases large enough to justify initial investment. The goal is not completeness but credibility, by creating early proof that sovereign environments can operate reliably, securely and at scale.
- **The Second Wave:** Concentrates on scaling shared infrastructure and data ecosystems. With demand signals in place, ecosystems expand compute and energy capacity on bankable terms, industrialize operating models, and invest in sector-specific data products and data-sharing mechanisms. This is where many initiatives falter. Because they attempt to scale infrastructure without first resolving governance, operating model and talent constraints.
- **The Third Wave:** Builds durable advantage and exportable capability. Ecosystems deepen specialization in selected domains, support a competitive provider landscape, and enable start-ups and integrators to scale. At this stage, trusted capabilities become not just domestic enablers but sources of regional or global differentiation.

The most common failure mode is mis-sequencing. This occurs when there's heavy investment in shared assets before demand and governance are ready. Or by pursuing global leadership ambitions without the data, adoption, and operating foundations required to sustain them.

"Ultimately, sovereign AI is not about full-stack independence; it is an ecosystem play," McKinsey says. "Those who orchestrate coherent systems—in which sovereignty is applied at critical control points, and governments, solution providers, and enterprises align incentives—could turn infrastructure into trusted capabilities and turn trusted capabilities into scaled outcomes."

Since we started with one C-Suite fable, let's end with another: The CFO unveiled the company's new AI sovereignty validation framework at the board's strategy meeting. "This framework ensures full financial transparency, algorithmic accountability and zero dependency on external black-box services," he said, scrolling through a deck bursting with acronyms. "The sooner we deploy, the sooner we reduce operational risk."

The board murmured approvingly, until the vice chair leaned forward. "I like it, especially the part about the AI having complete operational autonomy," he smirked. "Just make sure every autonomous action comes with a cost-benefit analysis, a three-scenario forecast, and a footnote explaining why we shouldn't panic."

Raju Chellam is a former Editor of Dataquest and is currently based in Singapore, where he is the Editor-in-Chief of the AI Ethics & Governance Body of Knowledge, and Chair of Cloud & Data Standards.
maildqindia@cybermedia.co.in