

Building Al value within borders

Accenture and NVIDIA accelerate nation-building AI solutions with the "Operating System for Sovereign AI Clouds"

The global AI imperative

The AI imperative is gaining momentum not just for businesses, but for the countries and regions in which they operate. Global leaders recognize Al's rising potential to impact their economies and their GDPs. And as AI permeates every facet of industry and society, there is a rising need for local ownership and control of all elements of the AI stack-from infrastructure to models and applications. Those that consider data and AI sovereignty as crucial to their long-term strategies are investing heavily in AI across key sectors in their countries, creating new kinds of value. Those who opt out of sovereign AI investment may forfeit future relevance on the global stage.

Cop

Why Sovereign AI?

Sovereign AI is an extension of the Sovereign Cloud movement that began in Europe several years ago, in the wake of regulatory frameworks targeting data usage and protection. The Sovereign Cloud movement is now seen as a key driver of economic and national security and largely focuses on creating environments that enable organizations to safeguard and control their data within borders, ensuring ongoing security and compliance. Sovereign clouds give governments and enterprises control over their own infrastructure, helping them guard against the possibility of unplanned disruptions brought about by geopolitical or economic shifts.

When it comes to Sovereign AI, the same drivers stand, but the stakes have risen. It's no longer *just* about control and data protection—it's an opportunity to reinvent how value is created for your region or country. Organizations are waking up to the fact that they don't want to outsource their local data and intelligence, only to buy it back at a premium from foreign service providers. AI adoption is also driving deeper concerns about the handling of a nation's core data assets across private and public sectors. Countries need to build resilience and autonomy into their national security and foster an AI economy that powers their national and economic security.

Early adopters of Sovereign AI include telecommunication providers, governments, NGOs, research institutes and enterprises in strategic industries. However, today, Sovereign AI is also scaling beyond this group to benefit a broad range of organizations across the globe. Sovereign AI adopters are already realizing opportunities to advance AI solutions that benefit their countries and citizens, while protecting their data and interests.

What is Sovereign AI?

Sovereign AI is a nation's ability to develop and deploy AI capabilities, leveraging its own infrastructure, data, and talent to foster innovation, drive economic growth, and advance its strategic interests—while maintaining control over the entire AI lifecycle.

How sovereignty drivers have evolved in the era of AI

Control and compliance

- National AI strategies and policies governing AI usage are beginning to crop up globally and are growing rapidly.
- At least 69 countries have proposed over 1,000 Alrelated policy initiatives and legal frameworks to address safety and governance.¹
- Adversarial AI and hardware supply chain attacks are creating an entirely new market for AI governance, security and compliance.

Technology independence

- As AI is integrated into critical infrastructure and services—like hospitals, financial institutions and public services—the risk of relying on services delivered from out-ofcountry providers grows.
- Nations want to decrease dependency on other countries for their own security and intelligence needs.
- Encouraging the adoption of locally developed AI tools and platforms, along with leveraging open source generative AI models, can accelerate the growth of national AI economy.

Protecting culture and values

- The ability to deliver localized models that represent local language, culture and priorities—and act accordingly—is a game-changer for many regions.
- The continued democratization of AI is accelerating a need for more localized models: E.g., if the new programming language is natural language, models that support local languages will be critical to local technology development.

Who can benefit?

The economic potential of AI, coupled with a shift in focus from risk management to value creation, has made sovereignty a key priority in every country. A new set of national champions will emerge to power their Sovereign AI economies, ultimately altering the global ecosystem. And the race is already on. Recently, at GTC 2025, NVIDIA CEO, Jensen Huang, asserted that global expenditures on data center buildouts will reach \$1 trillion by 2028.² NVIDIA is also projecting \$10B in near-term revenue from government-endorsed Sovereign AI investments—up from zero projected last year.³

Governments employ up to 30% of the population of any given country,⁴ and they face demographic challenges, from aging population to talent shortages. Globally, the ratio of working-age individuals to those aged 65 and up is expected to shrink from 6.4 in 2024 to 3.9 by 2050.⁵

These challenges can be solved with the adoption of sovereign AI systems that help governments unlock the efficiencies needed to serve their citizens, meeting required data sovereignty, computational localization and AI model localization.

Governments and enterprises who adopt sovereignty as a strategic imperative will create significant value as they become AI leaders in their respective regions. They can quickly monetize investments in country-specific solutions and advance their global leadership by providing nation-building services in areas like healthcare, banking, smart cities and other key innovation sectors. Some groups, which we call "Sovereign Frontliners"—usually nationally endorsed entities—stand to benefit most from Sovereign AI investments:

Sovereign Frontliner groups

Telcos and infrastructure providers

These companies are often the frontliners to implement Sovereign Al, as trusted national technology providers with access to data centers, connectivity and reach. They stand to monetize infrastructure investments, attract regulated industries and offer connectivity for core use cases, including those where proximity is key. Indonesian telco, Indosat Ooredoo Hutchison, is a prime example of a frontliner leading the charge in creating national AI capabilities.

Government/ national issue-driven organizations

Government organizations often invest in domestic infrastructure and influence sovereign agendas to advance their own economies. For instance, HUMAIN, a subsidiary of The Kingdom of Saudi Arabia's Public Investment Fund focused on AI, is driving the Kingdom's mission of becoming an international AI powerhouse with major investments in domestic infrastructure.⁶

Industry champions

Industry leaders can boost competitiveness, drive data collaboration and control compliance across their value chains with AI solutions tailored to industry-specific needs. For example, Italian aerospace, defense and security company, Leonardo, is actively investing in supercomputing, cloud and AI to drive innovation within their ecosystem.⁷

accenture | 📀 NVIDIA

Accenture. All rights reserved.

The challenge

Moving from Sovereign Cloud to Sovereign Al raises new questions. Sovereign Al is built on many Sovereign Cloud principles, and expands sovereignty considerations across the development, training and deployment of sovereign and local-language models and Al applications.

Many organizations are looking for modular, composable stacks that they can combine to build Sovereign AI solutions. There are several technological options available, from hyperscaler-localized or dedicated sovereign AI clouds, to Neoclouds, with a focus on delivering scaled accelerated GPU infrastructure.

Many Sovereign Frontliners are also assembling the entire Sovereign AI stack from infrastructure through AI industry solutions (in a colocation site or on-premise), often endorsed by governments. A prime example is the Sovereign AI deployment from Indonesian telecom, Indosat Ooredoo Hutchison (IOH), powered by Accenture, NVIDIA and Rafay (see **Client Story** below). Frontliners assembling their own solutions face the unique challenge of delivering a fully sovereign stack while enabling a rapid-scaling and self-service experience on par with CSP offerings. To achieve this, they must deliver on four requirements:

- ✓ Scalability: Supports a diverse set of end users through a robust multi-tenant architecture spanning all layers of the stack
- ✓ Seamlessness: Solves for the developer/operator experience of managing a large infrastructure fleet, allowing for rapid consumption and management of Al solutions
- Relevance: Offers industry applications and/or development capabilities to create stickiness (vs. just GPU capacity)
- ✓ Enterprise-readiness: Embeds verifiable governance and attestations to enable scaling of regulatory-compliant agentic solutions across an enterprise.

Speed-to-market is critical to remain competitive and thrive amid change, which has become the new global norm. Engineering these sovereign platforms to support the needs of a country may add time and complexity. So how do you efficiently build and run a Sovereign AI Cloud that creates stickiness, offers a CSP-like user experiences and is both multi-tenant and secure? The solution: An Operating System for Sovereign Al Clouds

Accenture and NVIDIA—along with their broad ecosystem of partners including Rafay, ServiceNow, Dell and HPE—have developed the Operating System for Sovereign AI Clouds. This holistic offering provides AI-ready infrastructure along with the capabilities to build and scale AI solutions-all within a sovereign environment. The Operating System includes Accenture AI Refinery, an AI platform that helps companies turn raw AI technology into customizable industry- and regionally-relevant solutions, co-developed and built on NVIDIA AI Enterprise software (see Figure 1). It incorporates both Accenture IP and NVIDIA NIM inference microservices, which provide secure and reliable AI deployment across a range of environments. AI Refinery also simplifies access to Al solutions and accelerates their adoption and scaling as part of a Sovereign AI approach.

The Operating System for Sovereign AI Clouds offers a solution that's highly flexible and controlled. It's rooted in Accenture's and NVIDIA's key strengths and existing partnerships to immediately allow a frontliner to go to market with AI solutions in their country. This offering and associated reference architecture address the six "layers" that you need to get right to run and externalize a sovereign AI cloud:

The Operating System for Sovereign AI Clouds

Applications	 Design, build and deploy applications spanning classical, generative and agentic Al to power up business and research. ✓ Enables the production and consumption of industry-specific agents with tailored, verifiable standards and compliance. ✓ Al Refinery provides tools for designing, testing and managing agents, compliant to sovereign policies. ✓ NVIDIA's agent Blueprints offer pre-built industry templates and governance frameworks to accelerate deployment. 	 Securely and responsibly operate and manage AI apps and services across the stack Efficient oversight across the AI stack—from infrastructure to applications—via a unified management interface. Control definition and enforcement points, secure LLM gateway and routing, comprehensive AI monitoring and FinOps using AI Refinery's AI control plane Shared services for management, maintenance and support. Self-healing systems and service that enable near real-time response to events including availability and service delivery. Automated compliance reporting and governance across data, consumption, hardware allocation and deployments.
Model	 Build, train and deploy classical and generative AI/ML models for use cases across sovereign business and research domains. Drives the adaptation of models to local, language, laws and cultures while storing them locally. AI Refinery enables highly customizable training with localized knowledge, without risking sensitive assets. Optimizes both training and serving of AI models within sovereign boundaries to ensure safety and quality. 	
nnowieage	 Harness and manage data securely across usage domains and ensure it's in an AI-ready and consumable form. Simplifies multi-modal knowledge extraction, curation and serving with Prebuilt tooling from AI Refinery, making enterprise data accessible for AI/ML. Secure and efficient processing of structured and unstructured data to extract precise knowledge for sovereign AI usage, enhancing performance and minimizing risks. Comprehensive knowledge repository, compliant with relevant usage policies. 	
Urcnestration	 Provide the infrastructure elasticity, scalability, automation and reliability needed for high efficiency. Rafay's central orchestration platform facilitates efficient, self-service infrastructure and Al application management. Seamless front-end tenant console, delivered through ServiceNow (or partner of choice), enhances user experiences and simplifies the adoption of Al solutions. Accenture's specific security and automation controls facilitate seamless, secure Al compute deployments. 	
	 Scale AI compute, storage and network capabilities for supporting broad sovereign usage. Ensures performance, security and compliance to support scalable AI workloads. NVIDIA GPUs supply the accelerated computing for complex AI models and tasks. Addresses multi-tenancy and sovereignty requirements across deployment and management. 	

06 Operations & Governance

accenture | 💩 NVIDIA.

s &

04

03

02

Figure 1: Accenture AI Refinery, built on NVIDIA: Turning raw agentic technology into useful business solutions



Copyright @ 2025 Accenture. All rights reserved. Accenture Confidential Information and IP - Not to be used or shared without Accenture's consent.

What we deliver

So, what does it mean to have an Operating System for Sovereign AI Clouds that functions at the level of, and often alongside, CSP solutions?

From the **owner/operator side**, it means gaining the ability to scale fast and efficiently plan for growth. It enables true multi-tenancy to support a wide range of consumption patterns, and a flexible, scalable, resilient architecture that can meet varying customer needs.

In turn, **customers and tenants** get a seamless experience where they can deploy services with a single touch—taking weeks of work down to minutes. They also benefit from pre-built integration patterns and embedded security and compliance that reduce complexities and accelerate value. Frontliners who embrace the Operating System for Sovereign AI Clouds aren't competing with CSPs. But they are offering a trusted, seamless and sovereign alternative for customers who are looking to maximize the value they can bring to their industries and regions. The Operating System for Sovereign AI Clouds is designed to support hybrid deployments, enabling seamless integration between hyperscaler-localized services and customer-built sovereign stacks. Its modular architecture and pre-built connectors allow providers to orchestrate workloads across environments while enforcing sovereign controls on data, models and agents. This ensures flexibility, interoperability and faster time-to-value-without compromising sovereignty.

Client Story

Indonesian telecom company, Indosat Ooredoo Hutchison, drives the sovereign AI agenda for their country

Indosat Ooredoo Hutchison (IOH), through its subsidiary Lintasarta, has partnered with Accenture and NVIDIA in a strategic collaboration to accelerate its sovereign AI cloud platform, aimed at propelling AI-driven digital transformation in Indonesia.

IOH is developing applications for industry-specific use cases based on its new AI cloud, Sahabat-AI and the NVIDIA AI Enterprise software platform. Sahabat-AI is a collection of open-source Indonesian large language models (LLMs) that local industries, government agencies, universities and research centers can use to create generative AI applications. Built by Indonesians, for Indonesians, Sahabat-AI models understand local contexts and enable people to build generative AI services and applications in Bahasa Indonesian and various local languages.

The collaboration will initially focus on AI solutions for Indonesia's financial services sector, one of the key pillars of the country's economy. Powered by the Accenture AI Refinery platform, and Rafay's orchestration platform for AI Clouds, Accenture will provide pre-built solutions that can significantly reduce time to value, with a modular architecture to meet client needs wherever they are in their AI journey. These solutions will enable Indonesian banks to harness AI, driving profitability, operational efficiency and sustainable growth in a highly competitive market.

This collaboration underscores Indosat's larger purpose of empowering Indonesia, contributing to the nation's digital transformation and technological sovereignty. By developing the country's first sovereign AI platform, Lintasarta and partners are working to unlock the potential of AI while upholding the critical tenets of data security, sovereignty and innovation.

Foster resilience and fast-track growth

We're in an era of change and uncertainty. Those organizations who are ready to take on the sovereign agenda have an opportunity to represent their whole regions as AI competitors on the global stage. To do that, they need to invest in technology that allows for secure, localized and scalable AI development. They need to be ready to deliver value-based solutions, tailored to their local needs. And they need to move quickly.

The Operating System for Sovereign AI Clouds is designed to fast-track the deployment of AI solutions in a country, shifting the focus from internally engineering multi-tenant cloud features to enabling scaled, nation-building AI capabilities. It's grounded in several core partner strengths: Accenture's AI Refinery platform, coupled with deep industry expertise and specialized HPC infrastructure capabilities; NVIDIA's compute power and AI software offerings; Rafay's seamless orchestration platform; ServiceNow's experience-led software; and many other partner capabilities that provide the flexibility to meet individual customer needs.

By leveraging proven AI cloud patterns, Sovereign Frontliners can focus on driving consumption and paving a clear path to maximize both value and sovereignty in a fastmoving market, advancing the interests of their businesses and their regions

Ready to get started?

Connect with us today to let us assess your current stack, share examples of scaled sovereign AI solutions we're implementing with our customers, and map a path to building your own Operating System for Sovereign AI Clouds.

Visit accenture.com/contact-us

and connect with our contributors on LinkedIn.

Key Contributors

David Wood

Global Lead – Technology Consulting | Accenture

Mauro Capo Cloud First Lead, ICEG and Sovereign Cloud Lead | Accenture <u>https://www.linkedin.com/in/maurocapo/</u>

Kunal Shah

Managing Director – Strategy & Consulting, Growth Markets and Data & AI Lead, APAC | Accenture <u>in https://www.linkedin.com/in/kunalushah/</u>

Atish Ray

Chief AI Architect | Accenture

Bryan Rich Global AI Lead – Health, Public Services and Defense | Accenture <u>https://www.linkedin.com/in/bryan-rich/</u>

Justin Stark

Chief Architect – Sovereign Cloud, APAC | Accenture

Chris Penrose

Global VP of Business Development – Telco | NVIDIA <u>https://www.linkedin.com/in/chris-penrose/</u>

Waleed Badr Director, Global NVIDIA Cloud Partners – Telco | NVIDIA in linkedin.com/in/waleedbadr99/

Haseeb Budhani Co-founder and CEO | Rafay Systems in linkedin.com/in/budhani/

accenture | 🐼 חעוטת.

References

- 1. https://www.mindfoundry.ai/blog/ai-regulations-around-the-world
- 2. <u>https://blogs.nvidia.com/blog/nvidia-keynote-at-gtc-2025-ai-news-live-updates/</u>
- 3. <u>https://www.wsj.com/tech/ai/nvidias-new-sales-booster-the-global-push-for-domestic-ai-champions-6d005ab7</u>
- 4. <u>https://www.oecd.org/en/publications/2023/06/government-at-a-glance-2023_da193b0d/full-report/component-66.html</u>
- 5. <u>https://www.weforum.org/publications/global-economic-futures-productivity-in-2030/</u>
- 6. <u>https://nvidianews.nvidia.com/news/humain-and-nvidia-announce-</u> <u>strategic-partnership-to-build-ai-factories-of-the-future-in-saudi-arabia</u>
- 7. <u>https://www.ilsole24ore.com/art/leonardo-punta-leadership-supercalcolo-e-innovazione-AHFpvnM</u>

About Accenture

Accenture is a leading global professional services company that helps the world's leading businesses, governments and other organizations build their digital core, optimize their operations, accelerate revenue growth and enhance citizen services—creating tangible value at speed and scale. We are a talent and innovation led company with 750,000 people serving clients in more than 120 countries. Technology is at the core of change today, and we are one of the world's leaders in helping drive that change, with strong ecosystem relationships. We combine our strength in technology with unmatched industry experience, functional expertise and global delivery capability. We are uniquely able to deliver tangible outcomes because of our broad range of services, solutions and assets across Strategy & Consulting, Technology, Operations, Industry X and Accenture Song. These capabilities, together with our culture of shared success and commitment to creating 360° value, enable us to help our clients succeed and build trusted, lasting relationships. We measure our success by the 360° value we create for our clients, each other, our shareholders, partners and communities.

Visit us at www.accenture.com

About NVIDIA

NVIDIA is the pioneer of GPU-accelerated computing. We specialize in products and platforms for the large, growing markets of gaming, professional visualization, data center, and automotive. Our creations are loved by the most demanding computer users in the world – gamers, designers, and scientists. And our work is at the center of the most consequential mega-trends in technology.

Visit us at www.nvidia.com

Disclaimer: This content is provided for general information purposes only, it does not take into account the reader's specific circumstances or information made available after the date of publication, and is not intended to be used in place of consultation with our professional advisors. This document may contain AI generated content. Any reference to marks owned by third parties are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.

Copyright © 2025 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.