
Architecting Sovereign AI in the Middle East

Driving Resilience and Accelerating
National Growth

Architecting sovereign AI in the Middle East



Abir Habbal

Data & AI Lead, Middle East



Carla Maria Issa

Middle East Research Lead



Contents

Executive summary	4
What Sovereign AI means for the Middle East	6
Our Research	7
Infrastructure and data center sovereignty	8
Compute and inference sovereignty	9
Model sovereignty	9
Application sovereignty	10
Keeping value local	12
The five imperatives	
Imperative 1 – CEO Ownership: Make sovereign AI leadership’s new power play	17
Imperative 2 – Reframe sovereignty: From risk management to value creation	18
Imperative 3 – Expand your ecosystem: Navigate the sovereignty continuum	20
Imperative 4 – Redefine architecture: Sovereignty for intelligence	22
Imperative 5 – Reinvent Talent: Sovereignty for People	23
Conclusion	25
About the research	26
References & Acknowledgments	27



Executive summary

The Middle East has emerged as a strong performer on AI governance and sovereign cloud, but must now extend this momentum by building deeper sovereign capabilities across the full AI stack. This means retaining strategic agency across infrastructure, compute, data, models and applications at a time when AI systems underpin competitiveness and trust.

The Middle East is advancing AI sovereignty at a faster pace than many other regions. Our research shows

60% of business executives in the region describe their country's approach to AI sovereignty as "highly sovereign," with a strong national focus on local control of AI, data and infrastructure. Globally, only 32% say the same.

Organizations are acting on this ambition. Fifty-seven percent of data and workloads in the Middle East have moved to sovereign cloud environments, compared with 49% globally.

As geopolitical risks intensify, 65% of regional organizations report they are more likely to pursue sovereign solutions, slightly above the global average of 61%.



Globally, the race to scale AI has exposed tensions between speed and control. Enterprises and governments seek to capture productivity gains while ensuring regulatory compliance, data protection and resilience. In the Middle East, sovereignty is embedded in digital and cloud strategies, reinforced by centralized governance and strong regulation. The region leads in data control, and sovereign cloud adoption is advancing at the model layer and expanding infrastructure. Yet it remains integrated with global ecosystems at the most critical layers, compute and inference.


But ambition is outpacing capability, and dependence persists. This gap is strategic and widening.

Sixty-two percent of respondents in the Middle East believe local sovereign AI solutions lag global hyperscalers in capability and maturity. And 43% point to the limited availability of national AI solutions as a barrier to advancing AI sovereignty. As a result, only 35% say sovereignty efforts meaningfully reduce reliance on global providers.

As seen in broader digital transformation efforts, constraints in core technology capabilities directly limit an organization's ability to innovate, scale and respond to market shifts. In the context of AI, this dependence risks ceding competitive advantage to external ecosystems.

Building sovereign capability demands sustained capital, ecosystem co-ordination and deliberate trade-offs on dependency. More than three-quarters (77%) of respondents believe governments play a critical role in advancing digital and AI sovereignty.

The next phase will determine whether early momentum translates into a durable advantage. More than half of organizations in the region plan to increase investment in sovereignty-related capabilities in the next 12 to 24 months.



Based on our findings, five priorities stand out. First, sovereignty must become a CEO-led agenda, integrated with enterprise resilience, growth strategy and geopolitical risk management. Second, leaders should reframe sovereignty not only as a compliance or risk mitigation exercise, but as a platform for value creation and competitive differentiation. Third, organizations must design hybrid ecosystems that combine local trust and regulatory alignment with the scale and innovation capacity of global providers. Fourth, organizations should architect AI systems for flexibility across infrastructure, data and models to avoid locking themselves into a single technological path. Finally, organizations should move decisively from traditional workforce planning to a fully integrated, skills-based, human+AI talent strategy that aligns people, technology and business strategy from the outset.

What Sovereign AI means for the Middle East

The Middle East's recent advancements are due to its innate ability to simultaneously accelerate domestic growth while deepening international partnerships. Sovereign AI would enable technological choices that are flexible, secure and strategically governed, creating a fit-for-purpose AI stack that blends global and local components aligned with national priorities and risk thresholds.

With sovereign capabilities, the Middle East could selectively govern AI systems across data, infrastructure, compute, models and applications in ways that align with national priorities, cultural values and strategic resilience. It is not about full autonomy across the AI stack, but about securing control where it matters most.



Our Research

Our global survey of almost 2,000 business and government leaders across 28 countries found that geopolitical risk is accelerating the shift toward sovereignty: 61% are now more likely to seek sovereign technology solutions. In the Middle East, the intent is even stronger, with 65% of respondents reporting a greater likelihood of adopting sovereign approaches.*

Notably, 60% of Middle East respondents described their country's approach to AI sovereignty as 'highly sovereign,' reflecting a strong national focus on local control of AI, data and infrastructure. This stands in sharp contrast to just 32% globally.

However, ambition is running ahead of capability. While the Middle East leads on certain layers such as data sovereignty and regulation, structural dependencies persist, particularly around advanced chips, global software frameworks and high-intensity inference. The region is actively converting these constraints into catalysts through local investment, ecosystem development and targeted partnerships. The key challenge now is turning that momentum into durable, full-stack capability.

Artificial intelligence is rapidly becoming one of the most significant drivers of global economic value.

According to our research, adoption of generative AI alone could unlock more than \$10.3 trillion in additional economic value by 2038 for organizations if generative AI is adopted responsibly and at scale, industry by industry, value chain by value chain.¹

At the same time, AI-related capital expenditure is accelerating sharply, with analysts projecting AI investment could rise to around 1% of global GDP, underpinned by hundreds of billions of dollars in annual infrastructure spending by leading technology firms.

* Unless otherwise noted, all data points are from the Accenture AI Sovereignty Survey from our flagship research, ["Sovereign AI: Own your AI Future"](#), Accenture, 2025.

Understanding where sovereignty is won or lost requires examining the AI stack itself. The layers that underpin AI — compute, infrastructure, data, models and applications — each carry distinct dependencies, risk profiles and strategic implications. In the Middle East, progress across these layers is uneven: strong in governance and data, advancing at the model layer, and still structurally constrained at compute.



Infrastructure and data center sovereignty

At the infrastructure layer, sovereignty is more operationalized. Major hyperscalers such as Microsoft Azure, Amazon Web Services, and Google Cloud Platform now operate in-region facilities across Saudi Arabia and the UAE enabling local data residence, regulatory compliance and sovereign cloud configurations. Alongside these global players, the Middle East has developed a new class of local cloud providers purpose-built for national sovereignty. In the UAE, G42 Cloud anchors national data strategies with sovereign cloud services tailored to government, healthcare, and defense sectors.

In Saudi Arabia, SCCC Alibaba Cloud, a joint venture between Saudi Telecom Company and Alibaba Cloud offers in-kingdom cloud capacity aligned with data and regulatory mandates. Additional offerings from STC Cloud, and Mobily also reinforce infrastructure localization across sensitive industries.

More recently, HUMAIN has emerged as a significant force in Saudi Arabia's infrastructure sovereignty agenda, developing large-scale AI-optimized data centers anchored by strategic partnerships with NVIDIA² and AMD³ — signaling a deliberate push to build sovereign compute infrastructure capable of supporting frontier AI workloads at scale. These platforms reflect a broader shift that infrastructure sovereignty in the Middle East is being delivered through a blend of global partnerships and domestically anchored cloud ecosystems. In this layer, the Middle East is comparatively advanced.

Compute and inference sovereignty

Compute remains the most capital-intensive and geopolitically sensitive layer. While Saudi Arabia and the UAE are investing heavily in GPU capacity and AI-optimized data centers, advanced AI chips are designed by a small group of global firms and manufactured in concentrated regions. Saudi Arabia has secured orders for over 18,000 advanced Nvidia AI chips, signaling its intent to scale compute capacity rapidly. The UAE, meanwhile, has obtained licensed access to more than 20,000 Nvidia A100 and H-series chips, alongside earlier purchases of thousands of H100 units, reinforcing its ability to compete in large-scale model training and inference. Yet, raw chip acquisition does not guarantee full sovereignty at the inference layer. Even where data is stored locally and models are domestically developed, high-intensity inference workloads — particularly for generative AI — may still rely on globally hosted infrastructure due to cost, performance or capacity constraints.

Sovereignty at the inference layer is therefore inherently hybrid and likely to remain so in the medium term. Full compute sovereignty faces a structural ceiling: constrained by global semiconductor dynamics at the hardware level, and by infrastructure economics at the inference level. Together, these make compute the most critical bottleneck to long-term strategic autonomy.

Model sovereignty

The Middle East has made visible advances at the model layer. Saudi Arabia's ALLAM represents a domestically developed, Arabic-focused large language model aligned with national priorities. The UAE's Falcon and Jais demonstrate the ambition to anchor advanced and Arabic-optimized models within national ecosystems, supporting both global competitiveness and regional linguistic leadership. In Qatar, research-led initiatives such as Fanar, developed by the Qatar Computing Research Institute, reflect a deliberate focus on Arabic-language capability and sector-specific applications. These efforts signal progress toward model sovereignty, and even broader cultural reach, though large-scale proprietary LLM commercialization and hyperscale deployment remain earlier-stage relative to regional peers — an important distinction when assessing full model sovereignty.

Despite these advances, sustaining frontier model development requires continuous, large-scale investment to remain competitive — a challenge even the world's most well-resourced technology organizations face. Rather than treating reliance on global hardware ecosystems and foundational software frameworks solely as a constraint, the more strategic posture is to selectively leverage global model capabilities where they offer superior performance, while concentrating sovereign investment at the application layer — where local language, cultural context, regulatory alignment and sector-specific use cases create the greatest differentiation and value. This reflects a pragmatic and increasingly common approach as sovereignty does not require building everything domestically, but rather securing control at the points that matter most.


Application sovereignty

At the application layer, sovereignty becomes contextual and value-driven. AI systems deployed in government, financial services and critical infrastructure increasingly reflect local regulatory requirements, linguistic needs and cultural considerations. A growing ecosystem of regional software firms such as Presight AI and Mozn are playing a key role in embedding these priorities through localized AI applications, analytics platforms and domain-specific tools tailored to sectors like national security, finance and public services. These companies support application-layer sovereignty by building solutions aligned with regional regulations and national use cases.

Our research found that in the Middle East, 54% of organizations have prioritized sovereignty or residency requirements on the data layer first, followed by infrastructure at 31%. Just over half (59%) report 'high' sovereignty in 'Security and Governance,' reflecting early investment in regulatory compliance and risk frameworks. However, sovereignty across other layers remains partial with 52% of organizations indicating they have 'medium' or shared control across applications, data, AI models and infrastructure.

Figure 1: Application of Sovereignty Requirements

Q: Which of the following layers have you applied sovereignty or residency requirements on?



Sample size: Middle-East N = 185 (UAE = 85, Saudi Arabia = 85, Qatar = 15)

This points to a growing maturity curve. Organizations have largely secured the regulatory core of sovereignty but now must shift focus toward operational and technical control. For businesses, the message is clear: embedding sovereignty deeper across the AI stack is becoming a strategic lever for differentiation, resilience and innovation.

Taken together, the Middle East is constructing a layered and pragmatic sovereignty model strong in governance, security and data control, while rapidly building momentum at the model layer, including the development of Arabic-optimized LLMs. At the same time, the region is scaling infrastructure through a hybrid mix of hyperscalers and sovereign cloud solutions and mobilizing local software ecosystems to capture greater value at the application layer. Yet, despite this progress, structural dependencies persist particularly across global supply chains for compute and inference.

The strategic question is which layers matter most to secure national resilience, regulatory control and competitive edge. In this context, sovereignty is emerging as a targeted agency at the most consequential points in the AI value chain.



Keeping value local

This is where AI sovereignty becomes economically decisive. J.P. Morgan's analysis highlights not only the upside of AI, but also the scale of investment required to sustain it.

The firm estimates that the AI ecosystem would need to generate approximately \$650 billion in annual revenue by 2030 to deliver modest returns on the infrastructure currently being built.⁴

This level of capital intensity elevates AI from software discussion to a national infrastructure and industrial policy issue.

In this context, the Middle East has a distinct comparative advantage. It benefits from some of the lowest-cost and most stable energy supplies globally underpinned by abundant hydrocarbons, strategic investments in grid infrastructure and a growing base of renewable capacity. As energy becomes a key input in AI compute economics, the region's ability to power sovereign data centers at lower cost strengthens its position as a viable hub for high-performance AI infrastructure. This energy cost arbitrage not only enhances economic viability but also enables nations to retain greater control over where value is created and captured across the AI value chain.

Lower electricity costs reduce the total cost of ownership for sovereign data centers, making it more economically viable to host model training and inference workloads domestically. Over time, this advantage can help partially mitigate external dependencies by enabling local compute build-out at scale with greater economic autonomy.

AI sovereignty is a value-capture mechanism. Control over data, compute infrastructure, and AI platforms increasingly determines where productivity gains accrue, where high-value jobs are created, and which economies retain strategic leverage as AI scales. Jurisdictions that lack sovereign control over these layers risk becoming net importers of AI capability, even as they generate data and demand that fuel global platforms elsewhere.

For the Middle East, this framing is particularly salient. The region's early investment in Sovereign Cloud, data governance, and centralized AI leadership reflects an implicit recognition of these economic stakes. Rather than treating sovereignty as a constraint on innovation, Middle Eastern governments and enterprises are positioning it to retain strategic agency as AI adoption accelerates globally. Our research found that over one third (35%) of respondents in the region believe AI Sovereignty allows them to reduce their reliance on global AI providers.

Across the UAE, Saudi Arabia and Qatar, more than half of organizations are already migrating significant workloads into Sovereign Cloud environments. This places the region ahead of many global peers in execution, particularly for regulated and mission-critical data. Our research found that in the Middle East, 57% of data and workloads have already moved to Sovereign Cloud, compared with 49% globally.

Middle Eastern organizations recognize the potential value of Sovereign AI at both national and organizational levels. Half of respondents see significant value in sovereign AI platforms for national objectives such as security, data protection, and digital independence. At the enterprise level, value is associated with compliance, control, trust, and operational resilience.

However, this recognition has yet to translate into large-scale capital deployment at the organizational level. Adoption remains primarily compliance-driven, with 62% of respondents citing adherence to local regulation as the main driver—well above the global average. By contrast, only 35% identify reducing reliance on global AI providers as a priority, suggesting that sovereignty considerations are emerging but not yet driving broad-based investment decisions. The region's approach to Sovereign AI is best described as pragmatic rather than hesitant.




While organizations clearly recognize the strategic value of sovereign AI platforms, especially at a national and public-sector level — investment momentum is moderated by high infrastructure costs, limited availability of mature national solutions, and lingering concerns around performance parity with global hyperscalers.

As a result, adoption is proceeding selectively, with a strong concentration in sectors where regulatory alignment, national interest, and economic value clearly converge.

Importantly, 62% of respondents in the Middle East said local Sovereign AI solutions are ‘behind’ or ‘not competitive’ to solutions offered by global hyperscalers.

Figure 2: Challenges in Adopting Sovereign AI

Q: What are the key challenges your organization faces or anticipates facing in adopting or applying Sovereignty to AI?



Middle-East

Sample size: Middle-East N = 185 (UAE = 85, Saudi Arabia = 85, Qatar = 15)



This confidence gap does not reflect a lack of ambition. Rather, it highlights the distance that national platforms must still travel to meet enterprise expectations around performance, scalability, and ecosystem depth. Until this gap narrows, organizations will continue to balance sovereignty requirements with pragmatic reliance on global providers.

The race to create sovereign solutions provides opportunities for both global cloud providers and local suppliers to offer cloud-based infrastructure, data, cybersecurity and AI models. For companies, it creates a trusted stack that incentivizes regional industry and cross-industry collaboration and data monetization. For nations, sovereign AI also presents an opportunity to empower people and safeguard their rights, thereby enabling systems that reflect local language, culture and values. It gives citizens recourse if algorithms misstep and allows local institutions to capture more of the value. The reality is that sovereign AI turns AI from something imposed on people into something built for them.

On the enterprise side, respondents overwhelmingly see governments as a key catalyst:

77% in the Middle East say governments, states or institutions have a key role to play in increasing digital and AI sovereignty—through regulation, subsidies or investments.

Just over half (54%) of Middle Eastern organizations plan to increase investment in sovereignty-related capabilities over the next 12 to 24 months.



The imperatives

Companies seeking to capture the opportunities created by sovereign AI are moving fast, thinking big, and balancing offense with defense. Based on our research and work with leading organizations, five imperatives separate leaders from the rest.

Imperative¹

CEO Ownership: Make sovereign AI leadership's new power play

In the Middle East, AI sovereignty is not solely a technology consideration; it is integrated into national transformation programs, strategic industry development, and international technology partnerships. In this context, AI sovereignty is no longer solely an IT or innovation issue. It has become a strategic governance matter, as capital deployment, regulatory posture, infrastructure design, and foreign technology exposure increasingly shape long-term resilience and value creation.

Why it matters: Sovereign AI decides who has oversight over data, who shapes markets and who writes the rules of digital engagement. It influences trust, growth and geopolitics. That's why the strategy must start at the top. The more regulated the industry, the higher the stakes. In financial services, sovereignty determines who controls transaction data, trading models and customer identities—core assets that define competitive advantage and trust. In healthcare and life sciences, it governs access to patient records, genomic data, and clinical models—areas bound by strict regulation and national oversight. In defense and critical infrastructure, it dictates whether AI systems can safely operate on sovereign compute, free from foreign dependencies or interference.

Take action: Yet today, only 10% of organizations in the Middle East have made AI sovereignty a CEO or board-level concern. This is a critical blind spot. Without executive ownership, sovereignty risks becoming another compliance box to tick—rather than a lever of differentiation and value creation.

Organizations in the Middle East lead with 55% noting that a Chief Data or AI Officer is primarily responsible for AI sovereignty decisions, compared with 37% globally.

Concretely, this means AI sovereignty must move onto the CEO and board agenda as a core strategic priority. Leaders need to make deliberate choices about where control matters most, defining which data, systems and partnerships require sovereign oversight and where global providers remain appropriate. This requires clear accountability, with a single senior leader owning the mandate. And disciplined execution that embeds sovereignty into major investment decision—treating control and ownership as critical criteria when selecting platforms and partners.

AI sovereignty demands decisive action at the highest levels of the enterprise. Supplier choices are no longer purely technical decisions; they require balancing innovation speed with geopolitical risk, supply chain resilience, and long-term cost—tradeoffs that only CEOs and boards can adjudicate.

At the same time, organizations must coordinate across legal, financial, technological, and reputational domains to avoid fragmented execution and ensure governance frameworks align investment, partnerships, and compliance with both corporate and national priorities.

Finally, as sovereign AI architectures are actively shaped through policy and regulation, companies must engage proactively with policymakers and industry leaders to influence the rules of the game. Those that treat sovereignty as strategy will help define the AI era; those that do not will operate within terms set by others.

Imperative 2

Reframe sovereignty: From risk management to value creation

Our research finds that most organizations in the Middle East approach sovereign AI from a defensive posture — driven primarily by compliance requirements (62%), reducing reliance on global providers (35%), and maintaining control over critical data and models (23%). Only 7% cite monetization or cultural alignment as strategic motivators. In practice, sovereignty is being treated as a risk mitigation exercise rather than a value creation lever.

At the national level, however, a different model is emerging. Initiatives such as Saudi Arabia's ALLAM, the UAE's Jais, and Qatar's Fanar reflect a deliberate effort to build Arabic-language large language models tailored to regional linguistic, cultural, and regulatory contexts. These models are not simply about technological independence; they aim to enable localized AI deployment across government services, strategic sectors, and enterprise applications, strengthening digital ecosystems while expanding the addressable market for AI solutions.

Companies should flip the script. Sovereign AI can move beyond compliance and risk control to become a source of competitive differentiation — translating local trust, cultural alignment, and regulatory fluency into new products, services, and ecosystem advantages.

Why it matters: When companies build AI that speaks their market's language—literally and figuratively—they capture customer confidence and regulatory approval simultaneously. This trust unlocks collaboration that would otherwise be impossible.



Case Study:

From Compliance to Competitive Positioning — The UAE's Sovereign AI Play

Recent developments in the UAE illustrate an evolution in how AI sovereignty is being operationalized. Reports that OpenAI is in discussions with the UAE government and G42 to develop a culturally and legally tailored version of ChatGPT reflect a hybrid sovereignty model: pairing frontier global AI capabilities with localized linguistic nuance, regulatory alignment and embedded governance controls.⁵ Rather than building in isolation, the UAE is shaping how advanced AI systems operate within its national context.

At the same time, the launch of GovGPT, a generative AI platform designed for government use cases, demonstrates sovereignty embedded directly into institutional workflows.⁶ GovGPT supports public sector decision-making, policy development and service delivery within domestic legal and cultural frameworks. This represents sovereignty not only at the infrastructure or model layer, but at the application layer — where intelligence directly influences national operations.

Together, these initiatives signal a shift in posture. While many organizations in the region continue to prioritize sovereignty primarily for compliance or dependency mitigation, the UAE is demonstrating how sovereignty can be used to unlock value. Localized frontier models increase trust and relevance, accelerating adoption in regulated sectors. Government-embedded AI platforms improve institutional productivity and establish standards that catalyze broader ecosystem participation.

In this model, sovereignty becomes a mechanism for shaping markets rather than shielding from them — retaining strategic control while expanding economic opportunity within the global AI ecosystem.

Take action: Governments in the region play the leading role in turning sovereign ambition into economic and social impact. They can reinforce this momentum by funding local infrastructure, subsidizing trusted providers and opening public datasets in critical sectors to spark innovation for small and medium-sized

enterprises (SMEs), startups and research organizations. Partnering with academia is equally vital—building a pipeline of next-generation AI talent, developing new models locally and anchoring innovation at home. Providers, in turn, should act as ecosystem anchors, enabling smaller players to participate and thrive.

Imperative 3

Expand your ecosystem: Navigate the sovereignty continuum

By nature, the region is already operating in a hybrid environment. Building on national goals and vision plans, localization is happening at a rapid pace, but alongside critical international partnerships. Companies must determine where sovereignty truly matters, and how to dial it up or down depending on risk and value. Participating in sovereign markets and benefiting from national sovereign agendas doesn't require changing everything or making every workload sovereign. In fact, delicate geopolitical balancing has reduced concentration risk, while giving the region the best of US cloud partnerships, European financial regulation and Chinese hardware.

Our research shows that companies can meaningfully participate and capture sovereign opportunities by applying sovereignty measures to just one-third of their AI initiatives. Even within that third, the required degree of strict localization varies.

Why it matters: Not everything demands transformational change or heavy investment. Some use cases need only local data residency and oversight of global infrastructure, while others call for full-stack sovereignty—local data and infrastructure, multi-cloud and multi-model access, confidential computing, advanced encryption or even air-gapped systems when necessary. The future of AI infrastructure will be a continuum of choices that will allow organizations to fluidly combine global scale with local oversight to create fit-for-purpose solutions.

Take action: Organizations in the Middle East face distinct choices in selecting AI solution providers, shaped by national digital agendas, sovereign investment strategies and geopolitical considerations. Global hyperscalers (AWS, Microsoft, Google Cloud, Oracle) offer frontier innovation, scale and in-country cloud regions across the GCC. National champions and state-backed platforms play a central role, combining local trust, regulatory alignment and privileged access to strategic sectors. AI-native infrastructure providers ("neoclouds") complement this landscape by delivering GPU-dense, AI-optimized compute environments.



Case Study:

Take HUMAIN, for example. They have developed a broad ecosystem of partnerships spanning the full AI stack, underscoring its ambition to build sovereign capability from infrastructure to applications. At the foundation, strategic alliances with NVIDIA and Advanced Micro Devices Inc. (AMD) anchor large-scale AI compute expansion, including GPU-powered data centers and AI-optimized infrastructure designed to support model training and inference within the Kingdom.^{7,8} Collaboration with Qualcomm extends this architecture into hybrid edge-to-cloud deployments, strengthening distributed AI capabilities. At the cloud and platform layer, partnerships with Amazon Web Services enable scalable AI deployment while aligning with sovereign hosting and governance objectives.⁹



Federated alliances between governments, sovereign funds and global technology firms are becoming common, enabling countries to balance access to innovation with strategic oversight. To make this work, organizations must segment AI use cases by risk and value, determining where sovereign control is essential and where global platforms remain viable. They should then design a deliberate provider mix for each category, combining hyperscalers, national platforms and specialized providers as needed—while continuously reassessing this balance as geopolitical and regulatory conditions evolve.

Beyond infrastructure, HUMAIN has deepened its position within the global frontier AI ecosystem through strategic investment in xAI, linking domestic capacity with advanced model innovation.¹⁰ At the application and developer layer, collaborations with platforms such as Replit aim to stimulate local AI-native development and broaden enterprise and startup participation.¹¹ Collectively, these partnerships reflect a full-spectrum strategy: securing foundational compute, integrating global platforms, advancing model capabilities and catalyzing a domestic innovation ecosystem capable of building and monetizing AI across priority sectors.

Imperative 4

Redefine architecture: Sovereignty for intelligence

With business commitments and operating models in place, organizations should next assess their digital foundations against sovereignty requirements. This means aggressively evaluating hardware, data, cloud readiness, compute needs, cybersecurity and AI model selection for each critical use case.

Why it matters: Sovereignty now means more than just choosing where data is stored or managed and where the infrastructure sits. It also includes the intelligence itself—how models, agents and platforms talk, learn and make choices within set limits. This evolution demands a new class of architecture built for multi-cloud, multi-model and now multi-agent collaboration, where technologies share memory, exchange context and coordinate to solve complex workflows. A single workload might be drawing from multiple services and models from different providers with varying levels of sovereignty. They require error handling for probabilistic outcomes, compliance and guardrails embedded within every model and agent and governance that allows autonomy while preserving accountability, trust and resilience.

Take action: The first step is to conduct a sovereignty audit for each use case, considering factors such as country risk, industry regulation and business or use-case criticality. Based on this assessment, companies can determine the appropriate level of sovereignty required and apply technology accordingly. Then, they should benchmark their current infrastructure, data residency and provider dependencies to establish a clear baseline.

This benchmarking helps identify gaps and prioritize remediation efforts across the organization. As part of this process, firms must also detect systemic, interconnected threats woven into their digital core—issues that often span across applications, platforms and providers. Finally, organizations should establish dynamic controls and governance across the entire AI stack, enabling them to anticipate emerging risks and maintain systems that are resilient, trusted and defensible. This approach shifts sovereignty from a static compliance exercise to a living capability that evolves with regulation, technology and business needs.

Our research indicates that the Middle East demonstrates strongest sovereignty maturity in the data, cloud and governance layers. Thirty-four percent of organizations report full control over organizational data, with a further 55% controlling sensitive data, and 59% rate security and governance sovereignty as high. This reflects early investment in Sovereign Cloud adoption and centralized regulatory frameworks. Sovereignty at the model layer is advancing through regional initiatives, though large-scale training remains linked to global hardware ecosystems. At the application layer, most enterprises continue to rely on globally developed AI solutions, with localization occurring primarily through configuration rather than full domestic development.

Imperative 5

Reinvent talent: Sovereignty for people

With architecture and infrastructure foundations in place, organizations must turn to the human layer of the AI stack. Achieving Sovereign AI is not only a technical challenge — it is a talent imperative. This means moving decisively from traditional workforce planning to a fully integrated, skills-based, human + AI talent strategy that aligns people, technology and business strategy from the outset.

Why it matters: Sovereign AI cannot be built, operated or evolved without sovereign talent. Yet most organizations remain underprepared: only about one-third currently align their talent, technology and business strategies, and only a small portion of employees feel genuinely equipped with the AI skills their roles now demand. At the same time, AI spend continues to rise while investment in upskilling lags. This gap represents a structural vulnerability. True sovereignty requires more than access to AI tools — it demands the local human capability to build, govern and continuously improve them. That means treating talent not as a downstream consideration but as a foundational enabler of the entire AI stack, alongside compute, data and models.

Sovereign AI also demands a scalable ecosystem of talent, not just individual capability. AI-native teams, personalized career pathways and cultures of continuous learning and experimentation are the organizational conditions that allow human and machine intelligence to evolve together. Organizations that embed learning into daily workflows and redesign work around human-AI collaboration significantly outperform peers in adaptability, innovation and growth. For governments and enterprises alike, this extends to national talent pipelines, academia partnerships and workforce platforms that develop local AI expertise at scale.

Take action: The first step is to conduct a talent sovereignty audit — assessing current AI skill depth, identifying critical capability gaps across the AI stack and mapping dependency on foreign expertise or globally managed talent platforms. Based on this, organizations should redesign workforce strategy around dynamic skills visibility and continuous development rather than static roles, embedding AI-powered learning directly into the flow of work.



Leaders should invest in building AI-native teams with clear human-AI collaboration protocols and use AI to personalize career pathways and accelerate internal mobility. Simultaneously, governments and large enterprises should establish or deepen partnerships with universities, technical institutes and national workforce platforms to build sovereign talent pipelines at scale — ensuring a sustained supply of local expertise in AI development, governance and operations.



Finally, organizations should embed talent reinvention into leadership accountability. Senior leaders must model AI fluency, champion cultures of curiosity and experimentation, and treat continuous learning as a strategic system — actively measured, managed and adapted. This shifts talent from a support function to a sovereign capability in its own right: one that evolves in step with the technology it is built to govern.

Our research highlights both the urgency and the gap. Sixty percent of Middle East organizations expect to increase their sovereignty investments over the next 12-24 months yet lack of access to Sovereign AI talent and domain-specific expertise ranks among the top challenges organizations face in adopting Sovereign AI. At the same time, most respondents acknowledge their country's AI sovereignty approach remains a work in progress — capable but still dependent on external technologies and standards. Together, these trends create a compounding risk: investment is accelerating, but the local human capability needed to build, operate and govern sovereign AI systems is not keeping pace. For governments and enterprises alike, closing this gap requires treating talent not as a downstream workforce planning concern, but as a foundational sovereignty priority — investing in national pipelines, academia partnerships and AI-native workforce platforms that ensure the region can not only access AI, but own and evolve it independently.

Conclusion

The Middle East's experience suggests that Sovereign AI maturity will be achieved through deliberate, disciplined execution. The region has already established a strong foundation through sovereign cloud adoption, data governance and centralized accountability. The next phase will depend on closing capability gaps by improving the competitiveness of national AI platforms, reducing infrastructure costs and demonstrating measurable business and national value to power a sovereign AI economy.

Strategic use of energy advantages will also be critical. By capitalizing on some of the world's most cost-effective and secure power supplies, Middle Eastern nations structurally lower the cost of scaling sovereign compute. This not only reduces operational costs but also positions the region as a global hub for training and deploying advanced AI systems.

This creates a narrow but critical window. The Middle East is uniquely positioned to convert its governance-led approach into sustained competitive advantage, but only if it moves beyond control toward capability and scale.

The direction is clear. Sovereignty must extend across the AI stack, from infrastructure to models to applications, so that control translates into innovation, productivity, and economic value. The foundations are in place. What happens next will determine the outcome: whether the Middle East shapes the AI economy on its own terms or operates within systems designed by others.



About the Research

This report is based on a [global study](#) that combined quantitative, qualitative and policy research to examine how governments and enterprises are advancing Sovereign AI and Sovereign Cloud. The global survey covered 1,928 organizations across 28 countries and 18 industries conducted between July and August 2025. For this report we exclusively looked at the data for Saudi Arabia, the United Arab Emirates and Qatar N=185.

Respondents were senior technology and policy leaders—CIOs, CTOs, Chief Data, AI and Risk Officers—from both public and private sectors. All represented organizations that consider sovereignty highly or critically important to their technology stack. The survey included 29 core questions on awareness, adoption, motivations, challenges and maturity across the digital core—data, cloud, AI, integration and governance.

Industries included: Automotive, Energy (Oil & Gas), Software & Platforms, Government and Public Sector, High Tech, Communications & Media, Insurance, Banking, Metals & Mining, Aerospace & Defense, Life Sciences, Industrial Equipment, Utilities, Capital Markets, Chemicals, Travel, and Health (excluding public healthcare).



References



- ¹ [Work, Workforce, Workers Reinvented in the age of generative AI](#). Accenture. 2024.
- ² [HUMAIN and NVIDIA Announce Strategic Partnership to Build AI Factories of the Future in Saudi Arabia](#). NVIDIA Newsroom. 2025.
- ³ [AMD and HUMAIN Form Strategic, \\$10B Collaboration to Advance Global AI](#). AMD Newsroom. 2025.
- ⁴ J.P. Morgan Chase & Co., "AI Capex — Financing the Investment Cycle" (November 2025). Internal research note.
- ⁵ [OpenAI in talks with Abu Dhabi's G42 to create specialized ChatGPT version for UAE](#). Semafor. 2026.
- ⁶ [Department of Government Enablement records landmark year under AI-native vision](#). WAM. 2026.
- ⁷ [HUMAIN and NVIDIA Announce Strategic Partnership to Build AI Factories of the Future in Saudi Arabia](#). NVIDIA Newsroom. 2025.
- ⁸ [AMD and HUMAIN Form Strategic, \\$10B Collaboration to Advance Global AI](#). AMD Newsroom. 2025.
- ⁹ [Amazon and Humain announce \\$5B strategic AI partnership in Saudi Arabia](#). Amazon News. 2025.
- ¹⁰ [Saudi Arabia's AI venture Humain invests \\$3bn in Elon Musk's xAI](#). Financial Times. 2025.
- ¹¹ [HUMAIN and Replit partner to build a Nation of AI Coders and accelerate AI innovation across the GCC](#). Zawya. 2025.

Acknowledgments

The authors would like to thank the following people for their contributions that made this report possible:

Surya Mukherjee, Gargi Chakrabarty, Kathleen Trickey, Amal Sebastian





About Accenture

Accenture is a leading solutions and services company that helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed across the enterprise, bringing together the talent of our approximately 786,000 people, our proprietary assets and platforms, and deep ecosystem relationships. Our strategy is to be the reinvention partner of choice for our clients and to be the most client-focused, AI-enabled, great place to work in the world. Through our Reinvention Services we bring together our capabilities across strategy, consulting, technology, operations, Song and Industry X with our deep industry expertise to create and deliver solutions and services for our clients. Our purpose is to deliver on the promise of technology and human ingenuity, and we measure our success by the 360° value we create for all our stakeholders. Visit us at www.accenture.com

About Accenture Research

Accenture Research shapes trends and creates data-driven thought leadership about the most pressing issues global organizations face. Combining innovative research techniques, such as data science led analysis, with a deep understanding of industry and technology, our team of 300 researchers in 20 countries publish hundreds of reports, articles and points of view every year. Our thought-provoking research developed with world leading organizations helps our clients embrace change, create value, and deliver on the power of technology and human ingenuity.

For more information, visit Accenture Research on www.accenture.com/research

Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date provided on the front page, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances, and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information.

Accenture does not provide legal, regulatory, audit, or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals.

This document refers to marks owned by third parties. All such third-party marks 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 © 2026 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.