POLICY OUTLOOK

TRT WORLD research centre

Türkiye and the Future of AI Sovereignty

Ravale Mohydin

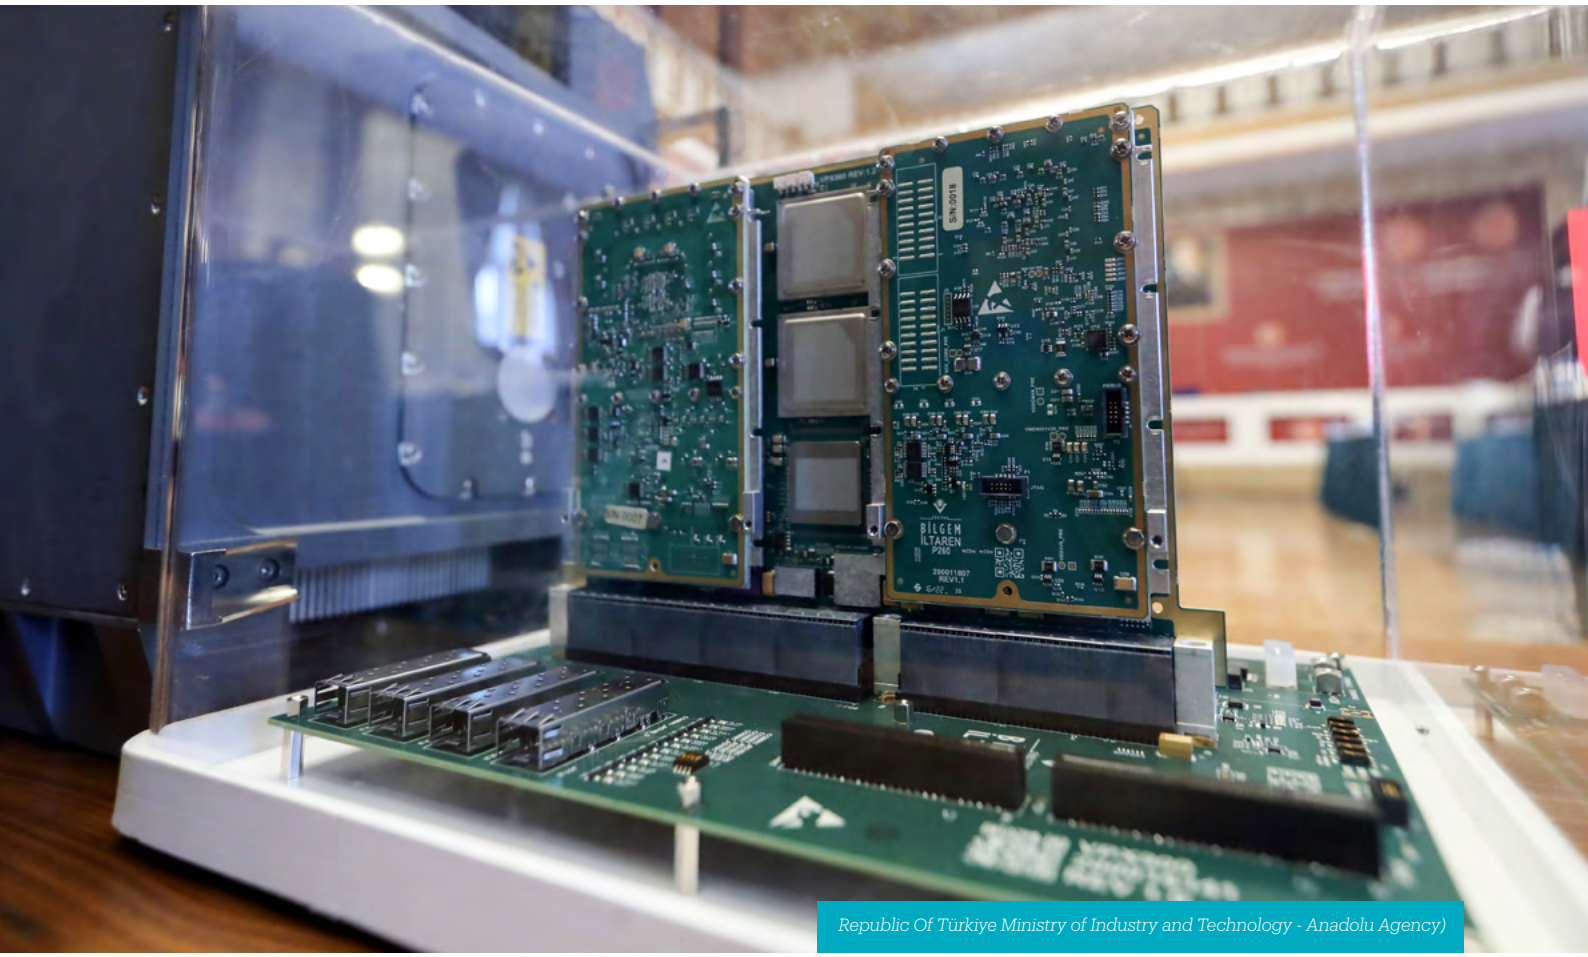

Republic Of Türkiye Ministry of Industry and Technology - Anadolu Agency)

As geopolitical competition progressively reshapes AI infrastructure, the question of technological sovereignty has moved decisively from policy aspiration to active strategic imperative. This policy outlook develops a six-layer stack model and seven-lever policy framework for evaluating sovereign AI capacity, applying both to Türkiye's structurally distinctive and increasingly credible position as a middle power with foundational sovereign capability already in place. Drawing on evidence from leading policy research institutions, it argues that durable sovereignty is achieved not through full-stack autarky but through managed interdependence, which is an approach Türkiye is better positioned to execute than most comparable economies.

Defining AI Sovereignty

The concept of AI sovereignty entered mainstream policy discourse around 2023 and has since been subject to considerable definitional contestation. In part, this reflects the genuine complexity of the underlying phenomenon, as well as the tendency of governments to deploy the language of sovereignty to signal strategic intent without committing to specific policy positions.

At its most fundamental level, AI sovereignty refers to a state's capacity to make independent decisions about the AI systems that shape its economy, security apparatus, as well as public institutions. In practice, this capacity operates across three distinct registers that do not necessarily move together. The first is infrastructural: control over the physical substrate on which AI systems run, encompassing data centres, semiconductor supply chains, power infrastructure, subsea cables and the cloud platforms that integrate these components. The second is operational: the ability to deploy, adapt and govern AI systems in sensitive domains, such as public finance, defence, healthcare and national security, without depending on opaque external platforms, foreign legal regimes such as the US CLOUD Act, or proprietary ecosystems that introduce switching costs precisely at the moment when political relationships become strained. The third is normative: the capacity to shape rules governing AI development and deployment at international standards bodies and governance forums, so that the architecture of global AI regulation reflects a state's values and interests rather than simply those of the actors who established the first institutional frameworks, i.e. some level of influence on the directions AI regulation takes worldwide.

It is important to note that these three registers are analytically distinct and do not co-vary in practice. For example, a country may achieve meaningful operational sovereignty by running sensitive workloads domestically on trusted infrastructure under national legal jurisdiction while remaining entirely dependent on foreign-designed chips and externally developed frontier models to power them. Conversely, a country may exercise significant normative influence in global AI governance forums while its domestic computing capacity is negligible by international standards. Understanding this disaggregation is the necessary precondition for developing a strategy that addresses the sovereignty challenge at the layers where it actually matters, rather than at the layers that are easiest to address.

Stack Layer	What It Comprises
Energy & Physical Infrastructure	Power generation, cooling systems, land, grid reliability and capacity
Semiconductors & Compute Hardware	GPUs, specialised AI accelerators, chip fabrication capacity and supply chains
Data Centres & Cloud Infrastructure	Colocation facilities, hyperscale cloud, emerging sovereign cloud region offerings
Data Assets	Training corpora, public datasets, sectoral and institutional data
Foundation Models	Pre-trained large language models and multimodal systems
Applications & Services	Sector-specific AI tools, public service systems, defence and industrial applications

State strategic planning must account for cross-stack layer dependencies, not just individual stack layer investments. For example, infrastructure without usable, well-governed data is strategically inert, and a domestically developed AI model without a trusted compute infrastructure cannot be deployed in sensitive domains where sovereignty considerations are most pressing. A sovereignty strategy that concentrates investment in a single layer without addressing cross-stack layer dependencies creates structural fragility rather than meaningful resilience.

The Case for Managed Interdependence

A working consensus has formed across the major research institutions, including Brookings, McKinsey, the Tony Blair Institute for Global Change and BCG, that full-stack autarky, which is a state of economic self-sufficiency where a country operates without any external dependence, including foreign aid or trading relations, is not actually a viable sovereign AI strategy for the overwhelming majority of countries. The semiconductor and foundation model layers, as noted above, are characterised by scale economies so extreme and supply chain concentrations so deep that attempting to replicate them domestically would require capital outlays and specialist talent pipelines that exist outside a small handful of actors. With that said, it is an argument for 'managed interdependence': a deliberate reconciliation of state autonomy with the international cooperation and commercial arrangements that are both necessary and, in many cases, strategically advantageous. In terms of actual operations, managed interdependence requires governments to develop classifications that segment AI systems by their sensitivity and strategic importance:

- Tier-one classification, which encompasses defence and intelligence systems, critical national infrastructure, as well as sensitive sovereign data assets requiring domestic compute and domestic legal jurisdiction as a baseline.
- Tier-two classifications, including sensitive public service delivery, financial system oversight and even national health data, tolerate sovereign cloud regions hosted by trusted foreign providers under national legal frameworks, provided data governance and auditability controls are in place.
- Tier-three classifications such as commercial applications, non-sensitive research and general-purpose productivity tools that can be run on global platforms without meaningful sovereignty risk.

Without this kind of explicit classification, sovereignty investment budgets may be applied without strategic logic, frequently concentrating on the most publicly visible or politically convenient layer rather than the most genuinely exposed one.

The urgency surrounding AI sovereignty stems from broader shifts, such as geopolitical power. Historically, states could remain dependent on external suppliers in discrete sectors without materially compromising their strategic autonomy. AI alters this calculation because it increasingly functions as a general-purpose technology embedded simultaneously within economic productivity, military capability, public administration and information ecosystems. As a result, dependence on external AI infrastructure is not merely a question of efficiency or procurement; it increasingly shapes a state's capacity to exercise independent policy choices. The sovereignty debate is therefore less about technological nationalism than about preserving strategic agency in a world where critical decision-making systems are becoming progressively more computationally mediated.

The Global Landscape

As of mid-2026, the sovereign AI agenda has moved decisively from policy aspiration to active capital deployment across a broad range of governments. A McKinsey survey of 300 executives, investors, and senior government officials, published in late 2025, found that 71 per cent characterised sovereign AI as either an existential concern or a strategic imperative for their goals. The scale of investment implied by this orientation is substantial: McKinsey estimates that sovereignty requirements will influence between 30 and 40 per cent of total global AI spending by 2030, representing a market of between 500 and 600 billion USD.

The United States-China Axis and Its Implications

The United States (US) and China occupy a structurally distinct position from all other actors in the global AI sovereignty landscape. Both countries have near-complete stack coverage: dominant semiconductor ecosystems, globally deployed hyperscale cloud infrastructure, widely used foundation models, and significant influence over international standards processes that are increasingly shaping AI governance norms. The intensifying rivalry between these two actors has become the organising force around which all other sovereignty strategies are being constructed. The successive rounds of US export controls on advanced AI chips, beginning with the Commerce Department actions of October 2022 and continuing through subsequent tightening measures, have forced a fundamental reassessment of supply chain assumptions across Asia, the Middle East and Europe. For every other country in the system, the central strategic challenge has become securing continued access to frontier AI capabilities without becoming structurally dependent on either pole of this rivalry in ways that compromise future policy autonomy. The table below showcases risks to national sovereignty at each of the stack layers described earlier.

Stack Layer	Sovereignty Risk Profile
Energy & Physical Infrastructure	Foundational to everything; consistently underweighted in policy discussions at the national level.
Semiconductors & Compute Hardware	Highest global concentration including Taiwan, the US, Netherlands as well as subject to active export control regimes.
Data Centres & Cloud Infrastructure	US hyperscalers control the majority of global capacity. Sovereign cloud options exist but are limited in scale.
Data Assets	Strategic layer where language, culture and institutional specificity create natural differentiation advantages.
Foundation Models	Extremely capital-intensive. Realistically, very few countries can compete at the frontier
Applications & Services	The most accessible layer for genuine domestic value capture and competitive differentiation.

The European Approach

Among democratic blocs, the European Union (EU) has pursued the most institutionally ambitious sovereign AI strategy. The majority of provisions in the EU AI Act, widely recognised as the world's first binding legal framework governing AI systems, will become applicable in August 2026, establishing a risk-based compliance regime with significant extraterritorial reach that is already shaping how AI developers globally structure their systems. Complementing this regulatory architecture, the AI Continent Action Plan commits approximately 200 billion EUR to AI infrastructure development, data centre capacity expansion and domestic industry support. At the same time, the October 2025 Apply AI Strategy targets the acceleration of AI adoption across the industrial and public sectors. These are serious financial commitments, and the EU's regulatory leadership at international governance forums is genuine. However, the structural constraint on European sovereignty remains largely unchanged: European data centres are substantially controlled by US hyperscalers, and the gap in foundation model capability relative to the US frontier reflects structural market concentration that budget commitments alone cannot overcome. In practice, European AI sovereignty is strongest in governance and norm-setting, and considerably weaker in compute infrastructure and frontier model development.

Differentiated Approaches Among Middle Powers

Middle powers are a term to describe countries that have placed themselves geopolitically between the US and China, without formally being committed to either, whilst having a substantial political economy themselves and having a "recognizable stake in a range of global issues", along with a capacity to exert influence on the outcomes of these issues. They are diverse in terms of geographies, population, economies and strategic cultures. These powers are likely to shape AI policies and governance models globally, if not already doing so.

Country / Bloc	Strategic Approach
Japan	**Alliance-based interdependence:** Tokyo has moved to underwrite outward FDI to build critical infrastructure nodes such as semiconductors, rare earth processing and 5G in trusted partner countries. The strategic objective is to reduce exposure to individual supply chain chokepoints through reciprocal access arrangements rather than full domestic replication.
UAE / Saudi Arabia	**Capital-as-sovereignty:** leveraging energy wealth and permissive regulatory environments to attract hyperscaler data centre investment at scale, creating leverage through hosting relationships that can be translated into governance influence and preferential access to frontier capabilities.
France / Germany	**Industrial sovereignty:** emphasis on domain-specific sovereign AI models for healthcare, defence and financial services, with continued resistance to full dependence on American platform infrastructure through the Gaia-X initiative and European cloud alternatives. Progress has been uneven, but the strategic direction is consistent.

The Compute Concentration Problem

A structural reality that sits beneath all national sovereignty calculations is the extraordinary concentration of AI compute infrastructure. Only approximately 30 countries globally currently host in-country compute infrastructure capable of supporting advanced AI workloads, and the nine leading public cloud providers, which are all either American or Chinese, represent roughly 70 per cent of the global AI compute market. Oxford Internet Institute research published in April 2026 further disaggregates this challenge by categorising three distinct levels of sovereignty when it comes to computing: the physical presence of data centres on national territory, the nationality of the companies that own and operate those data centres, as well as the nationality of the vendors (such as, for example, chip manufacturers) powering these AI systems. A country can perform well on the first measure, such as by establishing data centre presence on its territory, but still remain deeply dependent on the second and third measures. This disaggregation matters considerably for policy design, yet most national sovereignty discussions conflate these three levels, which is a major reason governments frequently make impressive-sounding data centre announcements that do not materially improve their actual sovereignty position. Policies also need to consider trade-offs; for example, a country can also introduce a policy to attract more data centres to its territory, but with that come social and environmental concerns related to the increased use of scarce resources, such as energy in the form of electricity, water, and large areas of land required.

A Policy Framework: Seven Strategic Levers

Drawing on evidence from McKinsey, Brookings, the Tony Blair Institute, BCG, as well as the World Economic Forum, the following framework identifies seven levers through which governments can expand genuine strategic agency in the age of AI:

Lever 1: Prioritise Before Committing to Infrastructure

As noted above, the most neglected step in national AI sovereignty strategy is, ironically, also the most foundational: determining with precision which classifications actually require sovereign treatment, because not everything can be justified within limited national budgets. Tier-one, tier-two, and tier-three classifications, as mentioned above, must be in place to ensure that sovereignty investment budgets are applied not just to the most visible classification but also to the most strategically exposed classification, and that genuinely critical systems are protected.

Lever 2: Build Domestic Capability Selectively and Deliberately

Rather than pursuing full-stack domestic replication, which is both capital-intensive and structurally fragile

at the layers where global scale economies are most concentrated, governments should identify 2 or 3 stack layers where domestic investment generates strategic returns that global alternatives cannot provide. Data assets represent the most consistent and defensible source of this differentiation, particularly where language, institutional knowledge, regulatory context or cultural specificity creates advantages that no external model trained on global corpora can replicate domestically. Domain-specific fine-tuning of foundation models and the development of sector-specific applications in defence, healthcare, and public finance sit in the same category. Full-stack ambitions, by contrast, tend to produce expensive fragmentation at precisely the layers where the gap with leading actors compounds rather than closes over time.

Lever 3: Build Regulatory Capacity as a Sovereignty Asset

States that adopt AI governance rules designed by others, either by regulatory default or as fast-follower compliance with externally developed frameworks, cede normative autonomy in a domain that will become increasingly consequential over the coming decade. Building genuine domestic regulatory expertise, engaging actively in International Organisation for Standardisation (ISO), International Telecommunication Union (ITU) and Organisation for Economic Co-operation and Development (OECD) AI standards development processes, as well as developing certification regimes that apply to both domestically developed and imported AI systems, constitutes a meaningful form of sovereignty in its own right. It is also considerably less capital-intensive than building domestic computing infrastructure.

A growing body of scholarship on what has been termed 'circular intelligence' offers a useful illustration of how this regulatory capacity argument operates in practice. Research examining the integration of AI into circular economy governance frameworks demonstrates that the most durable form of AI-enabled policy is not one in which AI systems are merely regulated from the outside, but one in which they are embedded within the governance architecture itself, used to synthesise policy across fragmented multi-level structures, to model the systemic consequences of regulatory choices before they are enacted and to enable what researchers describe as adaptive regulation, wherein policy instruments are updated in response to real-time environmental and economic data rather than fixed legislative cycles. The capabilities implicated in this approach, including natural language processing for policy synthesis across jurisdictions, machine learning for waste and material flow pattern detection, and digital twins for scenario testing across governance levels, are not frontier AI capabilities requiring cutting-edge compute. They are mid-stack applications that any country with a functioning data infrastructure and a modest domestic AI engineering base can develop and deploy. This is precisely why regulatory AI capacity constitutes a genuine sovereignty lever rather than merely a technical upgrade: it is achievable at a scale and cost that is realistic for middle powers, and it generates governance intelligence that is inherently non-transferable because it is trained on, and calibrated to, the specific institutional, legal, and environmental context of the country deploying it.

Lever 4: Use Public Procurement as a Sovereignty Instrument

Government AI procurement represents what is perhaps the most systematically underutilised sovereign AI lever available to most states. The potential of public procurement as an industrial policy tool to incentivise innovation has long been recognised by policymakers and researchers alike. By bundling public sector AI demand into multi-year framework agreements that embed sovereignty requirements such as data localisation obligations, encryption key ownership, system auditability and domestic staffing thresholds, governments create the demand signals necessary to justify private infrastructure investment at a scale that individual procurement decisions cannot achieve. This 'anchor tenant' model has been applied effectively in countries such as the United Arab Emirates (UAE) to catalyse significant investment in hyperscaler data centres. The mechanism is available to any government willing to aggregate public-sector demand rather than fragment it across agencies operating under independent procurement mandates.

Lever 5: Align Capital Instruments to the Stack's Risk Profile

Energy and data centre infrastructure requires patient, long-term capital with stable and predictable returns, a profile associated with infrastructure bonds, sovereign wealth fund co-investment, as well as development finance instruments. Foundation models and AI platforms require growth capital with significantly higher risk tolerance and longer investment horizons. Application development and enterprise adoption require venture capital pathways and procurement incentives. These are structurally distinct investment propositions that cannot be effectively served by a single public fund or a generalised AI investment programme. Sovereign AI ecosystems that are achieving results have designed distinct capital instruments for each stack layer rather than treating AI investment as a single asset class.

Lever 6: Invest in Talent at a Scale Commensurate with the Ambition

AI talent is another binding constraint that no infrastructure investment programme can substitute. The International Monetary Fund's (IMF) projects that effective AI adoption could contribute approximately 4 per cent to global GDP over the coming decade, and countries that lack the domestic human capital to deploy and adapt AI systems will see those gains captured elsewhere. Talent

development strategies for AI in particular must combine curriculum reform at secondary and tertiary levels, competitive international visa frameworks designed to attract specialist AI researchers and engineers, and sustained public-private co-development laboratory programmes that link academic training to real deployment environments. This last element is particularly important: the gap between academic AI training and the skills required for high-growth product development is a persistent structural problem that formal education programmes alone cannot resolve. Furthermore, the brain drain risk is acute for mid-income economies with strong technical universities like Türkiye, and this challenge is not amenable to resolution through strategy document commitments without a credible compensation and career proposition for qualified talent.

Lever 7: Shape International Governance Norms Early

For most middle powers, active, technically credible participation in multilateral AI governance is the most cost-effective sovereignty lever available. Engagement in processes including the G7 AI Code of Conduct, the Global Partnership on AI, the OECD AI Principles, as well as bilateral AI cooperation agreements with aligned states, allows countries to shape governance standards before they harden into de facto global requirements imposed by dominant actors. Countries that participate early and with technical substance gain disproportionate influence, commensurate with their relative technological capacity. Those that arrive late to these processes, as experience with earlier rounds of digital governance shows, inherit rules and frameworks built around other actors' interests and industrial structures. Clearly, effective AI sovereignty is not achieved through a single policy decision, a cloud contract, or a national model announcement. It is an ecosystem effort that requires simultaneous, coordinated action across these seven strategic levers. Countries that approach sovereignty as a coordinated design problem by asking how these levers reinforce each other will consistently outperform those that approach it as a procurement or communications exercise.

Türkiye: Strategic Position and Policy Trajectory

Now that a clear framework for benchmarking was delineated, one can say that Türkiye occupies a structurally distinctive position in the global AI sovereignty landscape, one that is neither straightforwardly advantageous nor straightforwardly constrained. It is not a tier-one power with full-stack technological dominance, and that cannot, arguably, be claimed by any country for now, but it is not a small or mid-income economy with no realistic options beyond dependency on external systems. As a G20 member, NATO ally, EU candidate state and a country that has demonstrated sustained and institutionally coherent capacity for defence-industrial indigenisation over more than two decades, Türkiye enters 2026 with the foundations, if not yet the full architecture, for a credible operational sovereignty strategy built on managed interdependence rather than on the strategically futile pursuit of self-sufficiency.

The Policy Architecture: Three Phases

Türkiye's AI policy development trajectory can be understood across three distinct phases, each building on the foundations established by its predecessor.

Phase	Label	Key Developments
2021–2023	Foundation	National AI Strategy 2021–2025 enacted via Presidential Circular 2021/18, making Türkiye the first country in Western Asia to establish a comprehensive national AI framework. The strategy established six strategic priorities, 24 objectives, and 119 measures, with coordination responsibility allocated between the Digital Transformation Office (CBDDO) and the Ministry of Industry and Technology. A national AI Steering Committee was also established.
2024–2025	Acceleration	The AI 2024–2025 Action Plan, published in July 2024, refined the strategic framework into 19 priority actions aligned with the 12th Development Plan (2024–2028). Ethical Guidelines for Reliable and Responsible AI were adopted in March 2023. The Defence Industry Sector Strategy 2023–2027 was published, and the draft AI Law (TBMM No. 2/2234) was introduced to parliament. A parallel bill (TBMM No. 2/3358) addressed deepfake regulation.
2026 onwards	Integration	The 2026 Presidential Annual Program represents a qualitative shift, embedding AI as a cross-cutting instrument of state capacity rather than a sectoral innovation tool. Two new directorates were created in December 2025: the Directorate General of National Technology and Artificial Intelligence, and the Public Artificial Intelligence Directorate General. The 2026–2030 National AI Strategy is in preparation.

AI as a Component of State Capacity

The 2026 Presidential Annual Program represents a qualitative shift in how Ankara conceptualises AI, and it is important to be precise about what is genuinely new in the document.

Previous presidential programmes acknowledged digital transformation and emerging technologies but largely treated AI as an auxiliary upgrade to existing systems, essentially something to be adopted selectively or piloted in limited domains without fundamentally altering how state institutions operate. The 2026 programme departs from this approach in a significant respect. AI appears throughout the document not as an innovation layer added to existing processes but as a cross-cutting capability embedded in the machinery of governance itself: machine-learning-based risk analysis in taxation and customs administration, automated decision-support systems in healthcare delivery, predictive analytics in social policy design, as well as continuous digital assistance for citizens through the e-Devlet platform.

This framing of AI as an instrument of state resilience and administrative continuity rather than of technological prestige carries genuine strategic significance for a country in Türkiye's position. It reflects an understanding, consistent with the analytical consensus among middle-power AI strategists, that the most durable form of AI sovereignty for a state at this level of development does not lie in owning a frontier model or in competing at the leading edge of global AI research. It lies, rather, in ensuring that state institutions can function effectively and without strategic vulnerability under conditions of economic stress, geopolitical pressure or administrative overload. The 2026 programme's emphasis on domestic AI models, secure data infrastructure and national compute capacity reflects a clear awareness of the risks associated with dependence on external systems in sensitive operational contexts, not a desire for isolation, but a determination to avoid the forms of dependency that could compromise policy choices or expose critical systems to external leverage later.

Defence AI: The Indigenisation Advantage

Türkiye's most advanced and institutionally embedded AI capability is in the defence sector, and the strategic significance of this for the broader sovereignty agenda is consistently underweighted in civilian AI policy discussions. The country has invested consistently and at scale over more than two decades in autonomous platform development, coordinated unmanned systems operations, advanced intelligence, surveillance and reconnaissance (ISR) capabilities and cognitive electronic warfare systems. The Steel Dome layered air defence concept and the Kızılelma aerial autonomy programme represent AI integration at the doctrinal level, shaping how the country conceives and executes military operations rather than just at the procurement or equipment level.

The significance of this defence-industrial AI programme for Türkiye's broader sovereignty posture extends beyond the military domain, as the potential for dual-use is being actively explored, with some military technologies migrating to civilian applications over time. The subsystems catalogued in the Teknokent Savunma Sanayii Kümelenmesi (TSSK) Exportable Products and Services Catalogue illustrate this point concretely: anti-jam GNSS receivers and controlled-reception-pattern antennas, developed originally for military aircraft and unmanned vehicles, are now being applied to civil-aviation ground stations, offshore energy platforms and autonomous mining operations; stabilised electro-optical/infrared gimbal payloads, originally developed for ISR and targeting, are being applied to wildfire monitoring, border surveillance, disaster search-and-rescue, and infrastructure inspection. The defence-industrial complex has been essentially functioning as a state-backed AI research and development programme with sovereign data as well as institutional knowledge, which is an advantage that purely commercial AI development strategies cannot replicate. Indigenised defence systems of the kind Türkiye has developed, often without explicitly framing them in sovereignty terms, necessarily generate domestic AI engineering talent, proprietary operational datasets, and institutional knowledge.

Infrastructure: The Gap Between Ambition and Current Capacity

Türkiye's current data centre capacity stands at approximately 250 MW, which represents less than 0.5 per cent of national electricity consumption. The 2030 Industry and Technology Strategy establishes a target of 1 GW of data centre capacity by 2030, supported by over 10 billion USD in projected investment. On the private sector side, initiatives such as the Nuvena sovereign AI infrastructure programme aim to establish domestic data centre capacity and a sovereign AI layer that would eliminate the need for Turkish startups and enterprises to host workloads and data outside the country. These are necessary and directionally sound commitments. With that said, the required energy dimension presents both constraints and opportunities. Türkiye's energy demand is projected to triple over the next three decades, and the additional loads associated with AI infrastructure and data centre operations are not yet fully captured in existing energy forecasts, as Energy Minister Alparslan Bayraktar has noted. At the same time, Türkiye's ongoing renewables buildout and its geographic position, which bridges European and Middle Eastern energy corridors, create conditions for positioning data centre growth as an infrastructure opportunity.

The semiconductor picture is more nuanced than a simple

absence of domestic capacity. Like virtually all states outside the US, China, the Netherlands and South Korea, Türkiye has no meaningful domestic chip fabrication capacity. Access to advanced AI accelerators, such as the Nvidia H100 and B-series hardware on which frontier AI development and deployment currently depend, is subject to US export control regimes that have been progressively tightened since 2022. Türkiye's approach, operating within the framework of the 2030 Industry and Technology Strategy and the Twelfth Development Plan (2024–2028), proceeds along three distinct lines. At the packaging and assembly layer, Türkiye is establishing Surface Mount Technology and advanced assembly infrastructure, which is a more tractable entry point than commercial-scale wafer fabrication and is already generating capability in smart-chip and defence electronics manufacturing, primarily in the Ankara technology zone cluster. At the design and development layer, national objectives as codified in both strategic documents place significant emphasis on indigenous semiconductor design and testing, with TÜBİTAK BİLGEM's UEKAE laboratories playing a central role in developing and validating chips for both defence and commercial applications. At the secure access layer, ASELSAN serves as the primary mechanism through which Türkiye maintains localised, sovereign access to the microelectronics its defence industrial base requires, managing supply relationships and developing substitution pathways that support strategic stockpiling and procurement security.

Yet the principal challenge facing Türkiye is unlikely to be a lack of strategic vision. Most of the major components of a sovereign AI ecosystem have already been identified in policy documents and development plans. The more difficult task lies in execution across institutional boundaries. AI sovereignty is inherently a coordination problem that requires energy planning, industrial policy, higher education reform, procurement strategy, regulatory development, and diplomatic engagement to move in parallel. Experience from both Türkiye and other middle powers suggests that the greatest implementation risks emerge not from insufficient ambition, but from fragmentation between institutions pursuing overlapping objectives without a sufficiently integrated governance mechanism.

Data Sovereignty and the Language Advantage

The 2024-2025 Action Plan identifies the development of a Turkish large language model as a high-priority deliverable, which represents strategically sound prioritisation and can be described idiomatically as "low hanging fruit". Language is the domain where global frontier models are most consistently and structurally weak in domestic contexts, and it is also the domain where Türkiye's institutional data assets encompassing legal, administrative, medical and financial corpora create genuine differentiation that no external model trained primarily on English-language or global multilingual datasets can replicate with the precision required for high-stakes public sector applications. A sovereign Turkish LLM, fine-tuned on public administration corpora and deployed across the state functions identified by the 2026 programme as AI-enabled, would provide precisely the form of operational sovereignty the strategy seeks to achieve.

In parallel, the Turkish government's planned Public Data Space (Kamu Veri Alanı), a structured and secure environment for researchers and startups to access public datasets, addresses a problem that is distinct from, but closely related to, data localisation. Open data portals already exist in the Turkish public administration. As McKinsey's analysis of sovereign AI ecosystems makes clear, keeping data within national borders does not automatically make it usable for AI development. For effective sovereignty at the data layer to take place, active data product development, interoperability standards that enable productive sharing across institutional boundaries, and sector-specific data consortia are required. The Public Data Space initiative aligns with this understanding, though its effective implementation will require sustained coordination among the public institutions that hold the relevant datasets.

It must be noted that the application case for a sovereign Turkish LLM extends considerably beyond public administration, translation, and citizen-facing services. The circular intelligence literature points to policy synthesis, including the automated aggregation and cross-referencing of regulatory instruments across ministries, municipalities, and international frameworks, as one of the highest-value near-term applications of NLP in governance contexts. For Türkiye, which operates a complex multi-level regulatory environment spanning national legislation, EU alignment obligations, municipal enforcement, and sector-specific frameworks in areas ranging from waste management to industrial energy use, this is a domain where a domestically-trained model with fluency in Turkish legal and administrative language would generate genuine institutional value that no commercially available multilingual model can replicate with the precision that consequential policy work requires. Furthermore, digital twin applications for scenario testing, enabling policymakers to model the distributional and environmental consequences of proposed regulatory changes before enactment, represent exactly the kind of sovereign application-layer tool that the Public Data Space, if properly constituted with real sectoral data, could support. These are not aspirational use cases. They are the kinds of applied AI governance tools that the 2026–2030 strategy's emphasis on AI as state capacity implicitly seeks, and naming them explicitly in the planning framework would sharpen both the investment logic and the procurement signals the government needs to send.

The Regulatory Framework

Türkiye's regulatory posture in the AI domain has been characterised by deliberate alignment with the EU AI Act's risk-based architecture, without full adoption of EU obligations, a positioning that reflects both the practical advantages of regulatory compatibility with Türkiye's largest trading partner and the desire to maintain national discretion over enforcement priorities and institutional design. The draft AI Law (TBMM No. 2/2234) proposes a risk-based regulatory regime with mandatory registration requirements for high-risk AI systems, turnover-based financial penalties, and a National AI Registry. A separate bill (TBMM No. 2/3358) addresses platform-specific concerns, introducing criminal liability for deepfake facilitation and a 6-hour takedown window for AI-generated harmful content, measures that align with Türkiye's broader approach to platform regulation in the digital domain. This dual-track regulatory approach is workable in principle, and it preserves a degree of national flexibility that full adoption of EU requirements would not.

The more pressing concern is timing. The draft AI Law (TBMM No. 2/2234) has been before parliament long enough that the regulatory uncertainty it creates has itself become a deterrent to the institutional and foreign investment required by the 2030 capacity targets. The TBMM AI Research Commission is expected to release a comprehensive sectoral report in late 2026, which may recommend additional legislation covering autonomous vehicles and AI in judicial proceedings. While this forward-looking legislative programme is very encouraging, the near-term priority appears to be resolving the uncertainty around the primary legislation that has already been drafted and is awaiting parliamentary approval.

The Diplomatic Dimension: Cognitive Diplomacy

Türkiye's most underexploited AI sovereignty lever is normative influence in international governance processes, and this is in some respects a more distinctive and durable advantage than the infrastructure investments that receive more policy attention. Türkiye has long been considered a successful practitioner of media and humanitarian diplomacy under the Justice and Development Party (AKP) government. Türkiye's institutional positioning, as a NATO member, EU candidate, G20 economy, and a state with deepening strategic relationships with Turkic states across Central Asia, African capitals and the broader Global South, provides it with legitimacy in AI governance discussions that no other middle power currently replicates. This legitimacy derives from the fact that Türkiye is not perceived as acting in the interest of any of the dominant technology blocs, which makes it a credible convener for the growing number of states that are seeking governance frameworks that are neither US-centric nor Chinese-centric in orientation.

Turkish researchers have articulated this opportunity through a Cognitive Diplomacy framework, which conceptualises diplomatic influence in the AI era as a function of three elements: presence in governance forums, demonstrable practice of responsible AI deployment, and resilience in compute and chip pipeline security. This can be a coherent framework for translating Türkiye's positional advantages into sustained governance influence. The strategic question is whether institutional machinery and diplomatic bandwidth exist to execute it consistently. The lessons from other domains of Turkish diplomatic engagement, as highlighted above, are that the positional advantage is real but requires sustained, technically credible follow-through to translate into durable influence.

Policy Recommendations

The following recommendations are structured by time horizon rather than by institutional ownership. This distinction matters because most AI strategy documents assign policy responsibilities without sequencing them, which leads to priorities being treated as simultaneously urgent rather than as ordered decisions that require different institutional efforts at different points in the planning cycle.

Immediate Priorities (2026)

- The draft AI Law requires enactment as the immediate legislative priority. The regulatory uncertainty generated by its continued draft status is already functioning as a deterrent to institutional and foreign investment, and no level of strategic ambition in other areas of the programme compensates for the absence of a clear legal framework governing AI development and deployment. Concurrently, the government should publish a workload classification methodology that defines with precision which AI systems require domestic compute infrastructure, which can be accommodated within sovereign cloud arrangements, and which can appropriately use global commercial platforms. Without this classification, infrastructure spending proceeds without strategic logic and the genuinely exposed workloads remain structurally unprotected.
- The Turkish LLM initiative requires a defined deployment target and a substantive public procurement commitment sufficient to justify the development investment. The language differentiation advantage that Türkiye holds in this domain is real, but it is not permanent, as other states and commercial actors are actively developing Turkish-language capabilities, and the window during which Türkiye can establish a leading position in this space is finite. Finally, a chip access strategy, whether structured through NATO procurement security arrangements, strategic stockpile development or

via investment in advanced chip packaging and assembly capabilities, needs to be formalised during 2026 if the 2030 compute capacity target is to remain achievable.

- On semiconductors, the immediate priority is not to initiate capability as Türkiye already has an Surface Mount Technology (SMT) and assembly infrastructure buildout underway, an active TÜBİTAK BİLGEM chip design and testing programme through UEKAE, and ASELSAN functioning as a de facto secure microelectronics access mechanism for the defence industrial base. What is absent is the formalisation of these three lines of effort into a coherent national chip strategy that connects them explicitly: one that establishes shared targets across the assembly, design, and secure access tracks; assigns capital commitments to each; defines the international partnership arrangements, (including NATO procurement channels, bilateral technology agreements or advanced packaging joint ventures) through which Türkiye will manage its continued dependence on foreign-fabricated advanced accelerators for AI workloads as well as creates a governance mechanism that prevents these three programmes from continuing to develop in institutional isolation from each other. This is a coordination and codification task, not a capability-building task from scratch, which makes it achievable within a 2026 timeframe if there is political will to convene the relevant institutions around a unified framework.

Medium-Term Priorities (2027–2028)

- The sovereign AI investment architecture requires redesign to align capital instruments with stack layers: infrastructure-grade capital for data centre development, growth capital for model development and platform investment, and venture pathways for application development and enterprise adoption. The current approach to treating AI investment as a single funding category results in systematic misallocation and leaves key layers consistently underfunded. In parallel, a TÜBİTAK-anchored defence-to-civilian AI transfer programme should be formalised to systematically commercialise the autonomy, ISR, and electronic warfare research that Türkiye has developed over the preceding two decades, targeting application in logistics, smart city systems, environmental monitoring, and public safety. The technology pipeline already exists within the defence industrial base; what is absent is the institutional mechanism to move it into civilian deployment.
- On the international front, Türkiye should formalise a Central Asian and MENA AI governance initiative under its own convening leadership, leveraging the Cognitive Diplomacy positioning to build normative influence at the multilateral tables where AI standards are being developed. Data corridor agreements with aligned partner states — enabling Turkish AI systems to access broader training data pools while maintaining domestic governance control over sensitive national datasets — would simultaneously strengthen the domestic model development programme and create durable bilateral relationships with states that currently lack credible AI governance partners.

Structural Priorities (2029–2030)

- Achieving the 1 GW data centre target requires a combination of public anchor investment to establish the demand signal and private sector buildout to achieve the scale, enabled by streamlined permitting processes and green energy co-location policy. Without the public anchor investment, private capital will not mobilise at the required pace or scale, and the 2030 target will remain aspirational rather than operational.
- On talent development, reaching 50,000 AI professionals with graduate-level specialisation requires not only curriculum reform and international talent visa frameworks but sustained public-private co-development laboratories that provide real deployment exposure rather than abstract training. The brain drain challenge is not addressable through policy incentives alone if the underlying career environment in terms of compensation, research infrastructure and career progression is not credible by international standards.
- The structural objective for 2030 is to position Türkiye as a tier-two AI governance norm-setter: not a rule-taker absorbing architectures and standards developed elsewhere and imposed through market power, but an active and technically credible contributor to the international standards processes that will govern the deployment of AI systems across the global economy for the following decade. This is a realistic ambition, but only if the diplomatic, technical, and institutional groundwork is consistently laid and followed through over the years.

Concluding Assessment

The global AI sovereignty agenda will intensify considerably over the next three to five years as geopolitical competition hardens supply chains, regulatory regimes proliferate and begin to conflict with each other and the gap between states with meaningful sovereign compute capability and those without widens into a structural divide that will be increasingly difficult to close. The IMF's projection of approximately 4 per cent global GDP uplift from effective AI adoption ensures that the economic stakes will sustain policy escalation regardless of the

broader geopolitical environment. Countries that lag in this transition will face not only slower economic growth but also diminished competitiveness across virtually every sector AI will reshape over the period.

For Türkiye specifically, the overall trajectory is optimistic. The 2026 policy architecture is more coherent, more institutionally grounded and more operationally specific than anything that preceded it. The dual directorate structure established in December 2025, the integration of AI into the Presidential Annual Programme as a cross-cutting state capacity instrument, and the emerging legislative framework collectively represent genuine policy maturation rather than rhetorical escalation. The defence AI programme, in particular, provides a sovereign data asset and engineering capability base (eventually translating into civilian products) that most comparable economies do not possess.

The critical uncertainties are concentrated on the supply side of the equation. Semiconductor access, the availability of long-horizon patient capital for infrastructure development and the retention of AI talent in a competitive global market for specialist skills are structural vulnerabilities that policy declarations and strategy documents cannot resolve in isolation. It must also be noted that the sequencing problem, i.e. the gap between ambition and sustained execution across political cycles and budget pressures, is not a new challenge for Turkish technology policy, or limited to it even, as strategic frameworks historically have encountered this as relatively straightforward tasks of articulating objectives turn into considerably more demanding tasks of maintaining coordinated implementation. The capacity exists, the institutional structures are being built, and the strategic intent is clear.

More broadly, the emerging AI order is likely to produce a new hierarchy among states, not based solely on technological sophistication but on their ability to manage strategic dependencies. For middle powers, the central challenge is not to replicate the capabilities of the United States or China across every layer of the stack. It is to identify the layers where dependence creates unacceptable strategic vulnerability and to concentrate resources there. In this sense, successful AI sovereignty in the coming decade will be measured less by technological self-sufficiency than by the capacity to exercise meaningful choice under conditions of interdependence.

The most credible and achievable strategic path for Türkiye is not full-stack technological autarky, which is a capital-intensive, structurally fragile and diplomatically isolating approach that would divert resources from the layers where genuine sovereign advantage is attainable, but the deliberate development of operational sovereignty at the layers where it matters most: defence, governance and critical public services. This is sovereignty conceived as strategic agency, defined as the capacity to make independent decisions and to shape the terms of engagement with global AI systems. It is the most achievable form of AI power available to the country in 2026, and, if built with coherence and sustained over the medium term, one of the most durable.